

Orchestration of Quantum Networks

Aliro



Orchestration of Quantum Networks

Summary	1
Introduction	1
Entanglement-based Quantum Network Orchestration	1
Evolution of Network Management.....	1
Modern Network Management	2
Two approaches to network management.....	2
Addressing challenges.....	3
Deep Dive into Network Management Protocols and Data Models	5
NMDA.....	5
NETCONF and RESTCONF Protocols.....	6
YANG Data Modeling.....	7
TMN Layering Model and Management Systems.....	8
FCAPS Model.....	11
TMN Example Topology.....	12
TMN Example: Configuration.....	13
TMN Example: Faults.....	14
Self-Healing Networks.....	14
Managing Entanglement-based Quantum Networks	15
Quantum network management vs. classical network management.....	15
Conclusion	17
References	19

Summary

Entanglement-based quantum networks represent a new frontier of technology. These networks leverage certain scientific principles, such as entanglement, to ensure unprecedented levels of security and performance. This white paper addresses the orchestration of such networks, exploring how they compare to classical networks.

Introduction

Entanglement-based Quantum Networks can support a wide array of applications, from secure government communications to financial transactions and critical infrastructure protection. In an evolving threat landscape, these networks are robust protection from sophisticated computational attacks now, and into the future when quantum computers are capable of breaking the security schemes we rely on today. Entanglement-based Quantum Networks utilize principles of quantum mechanics to secure communication channels with provable security and eavesdropper detection.

Orchestration and configuration of Quantum Networks is more complex than the classical networks we currently use for communications and transactions, but this complexity can be managed through principles originally developed for efficiently and effectively operating classical networks. In this white paper, we'll focus on relevant management principles of classical network orchestration on the orchestration of Quantum Networks, and also highlight how these systems differ and what the future holds for entanglement-based networks.

Quantum Network Orchestration

Evolution of Network Management

In the beginning, there was the Command Line Interface, or CLI. The command line was accessed via a serial port or a terminal server connected to a serial port, where operators would manually input commands. This method was labor-intensive and prone to errors, particularly due to the reliance on screen scraping. Screen scraping involved extracting data presented on the CLI screen, which was a cumbersome process often leading to tight coupling with human-oriented interfaces and command sequencing issues since CLIs typically have the need to have certain commands run before other commands in order for a procedure to be made effective. You can imagine how problematic this might be when considering the CLI presenting a table of data, and the script that's scraping the data has to process that tabular data in order to effectively extract the data from it.

With the advent of TCP/IP, network management began to evolve. The introduction of TCP allowed for remote access and the execution of CLI commands over the network, significantly improving efficiency. However, this new capability also brought its own set of

challenges. The initial implementations were fraught with issues such as the need for programmatic APIs and the difficulties associated with screen scraping, which persisted despite the improved connectivity. Even with a library of scripts, the scripts themselves needed to be sequenced manually. These early network management systems required heavy investments in scripting and manual interventions, underscoring the need for more automated solutions.

In response to these challenges, the Internet Engineering Task Force (IETF) began to develop standards to address the growing complexities of network management. In the late 1980s and early 1990s, protocols like Simple Network Management Protocol (SNMP) were introduced. SNMP became the recommended standard for network management, providing a framework for monitoring and managing network devices. Despite its widespread adoption, SNMP had limitations, such as its binary nature, which made it difficult to debug and express data models effectively. These shortcomings led to continued reliance on proprietary solutions and scripting.

The turn of the millennium ushered in significant advancements in network management practices, leading to modern tools for network management. The introduction of the Telecommunications Management Network (TMN) layering model and the FCAPS model (Fault, Configuration, Accounting, Performance, and Security management) provided a structured approach to managing network resources. This period also marked the emergence of the Network Configuration Protocol (NETCONF) and the YANG data modeling language. These later innovations addressed many of the limitations of SNMP by enabling more sophisticated and scalable management of network configurations and states.

Modern Network Management

Two approaches to network management

In modern network management, there is a strong desire to use the network-is-the-master approach for managing configurations. In the network-is-the-master approach, each network device, such as routers and switches, autonomously manages its configuration and operational decisions. This method leverages the embedded intelligence within these devices, allowing them to make real-time decisions and adapt to changes locally. However, this autonomy introduces complexities in maintaining consistent configurations across the network, as each device may handle configurations differently based on vendor-specific implementations. The lack of centralized control can make it challenging to enforce uniform policies and troubleshoot issues effectively, as logs and state information are dispersed across multiple devices.

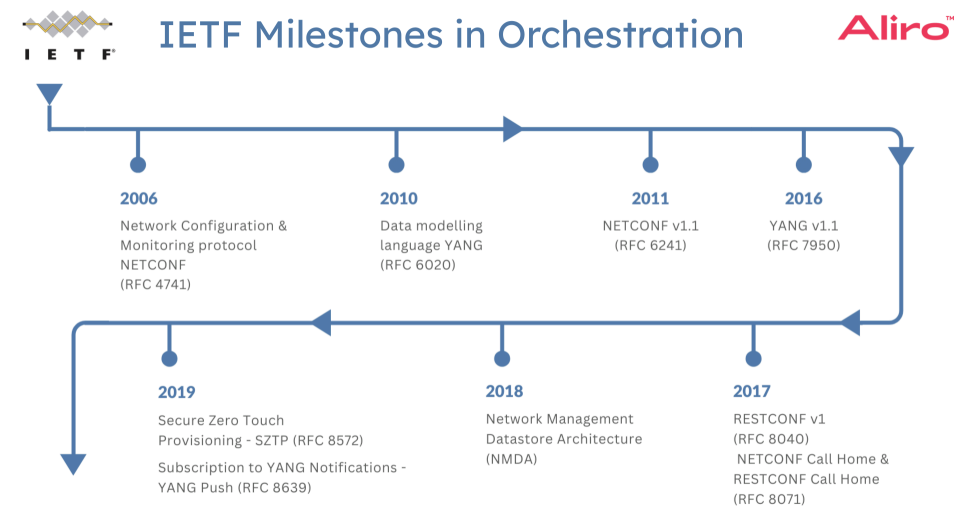
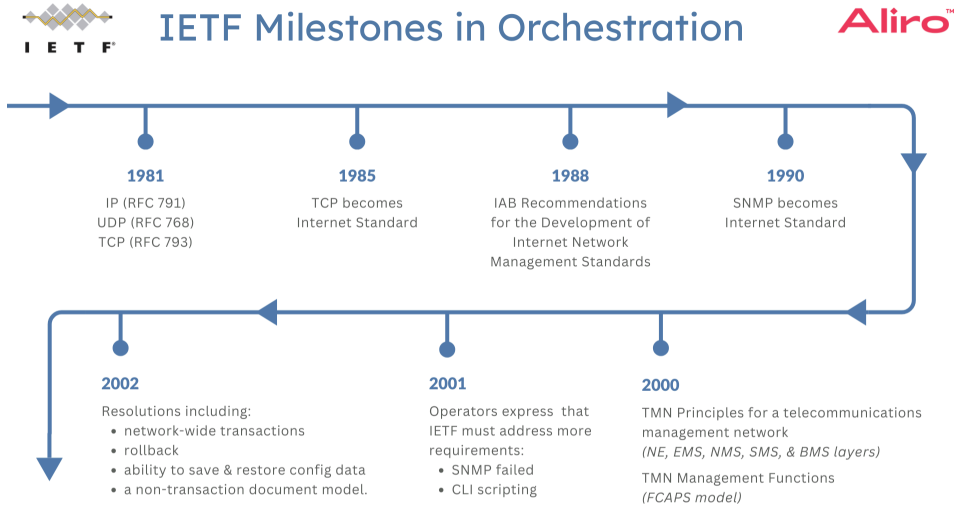
Scalability is another major issue with the network-is-the-master approach. As the network grows, coordinating configurations manually among numerous devices becomes increasingly cumbersome and error-prone. Each device's limited computational resources may also constrain the network's ability to handle complex orchestration tasks, resulting in potential bottlenecks. Furthermore, the diversity in vendor-specific protocols can create interoperability challenges, complicating the integration of new devices and technologies into the network. This approach might be suitable for smaller networks where device-level autonomy can be fully leveraged, but it struggles to maintain efficiency and consistency in larger, more complex environments.

“Management-is-the-master” is an alternative approach to the network-is-the-master approach. In this approach, the management system itself serves as the source of truth, centralizing all configuration data and pushing updates to the network devices. This approach, while potentially requiring robust synchronization protocols, provides a more controlled and consistent method for managing network configurations, ultimately enhancing the stability and reliability of the network.

However, the network-is-the-master approach isn't without challenges, particularly with overwriting out-of-band updates. Out-of-band updates are changes made directly on the devices, bypassing the central management system, which can lead to inconsistencies between the network devices and the management system. These discrepancies can create conflicts, increase the risk of errors, and complicate the synchronization process. In order to support the possibility of there being out-of-band updates, management systems should support the ability to synchronize those out-of-band updates or to import them into its data model as quickly as they occur.

Addressing challenges

The Internet Engineering Task Force (IETF) has played an important role in addressing the challenges associated with network management. As a key standards organization, the IETF is responsible for developing and promoting protocols that ensure the smooth operation and interoperability of the internet and related network technologies. Starting in the late 1980s and early 1990s, the IETF introduced protocols like the Simple Network Management Protocol (SNMP), which provided a framework for monitoring and managing network devices. While SNMP became widely adopted, it had limitations, particularly in terms of its binary nature and the difficulty of expressing complex data models. Recognizing these shortcomings, the IETF continued to evolve its standards, leading to the development of more advanced protocols and architectures that address the growing demands of modern network management.



Among the many significant contributions from the IETF are the introduction of protocols and architectures like NETCONF, YANG, and the Network Management Datastore Architecture (NMDA). NETCONF provides a standardized mechanism for installing, manipulating, and deleting the configuration of network devices, supporting transactional operations that enhance network stability. YANG, a data modeling language, enables the precise definition of configuration and state data, facilitating interoperability and automation. NMDA offers a structured approach to managing network data stores, ensuring consistency and reliability in network configurations. Together, these standards have transformed network management, enabling more efficient, scalable, and automated processes that meet the demands of contemporary digital infrastructures.

Deep Dive into Network Management Protocols and Data Models

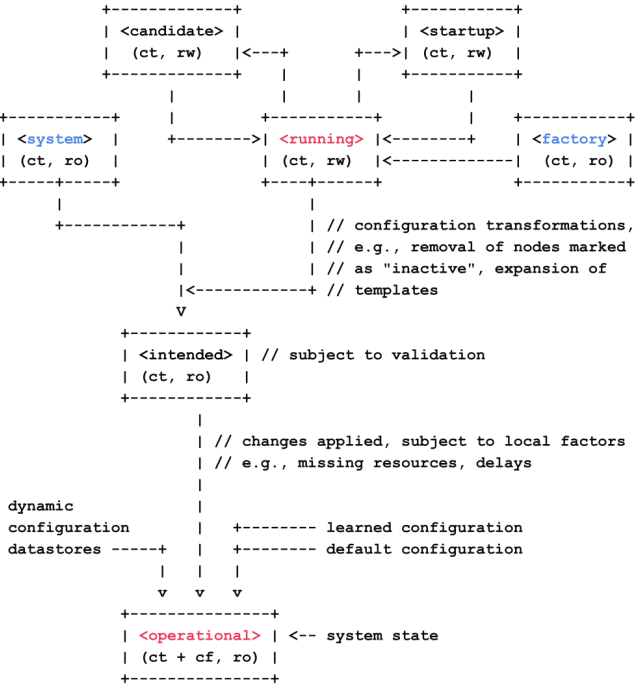
NMDA, NETCONF, and RESTCONF, supported by the YANG data modeling language, offer powerful capabilities for managing complex network configurations. NMDA provides a structured framework that enhances consistency and reliability in network management. NETCONF's robust transactional capabilities and RESTCONF's simplicity and flexibility make them complementary tools in a network operator's toolkit. YANG's clear distinction between configuration and operational state data enhances the precision and reliability of network management. Together, these protocols and data models form the backbone of contemporary network management systems, ensuring that networks can be managed efficiently, reliably, and at scale.

NMDA

Network Management Datastore Architecture (NMDA) provides a structured approach to managing configuration and state data across network devices. NMDA ensures consistency and reliability by defining clear roles for different types of datastores. A datastore is where all the configuration and operational state for the device resides. Each datastore serves a specific purpose, facilitating precise control over how configuration changes are made, validated, and applied.

Network Management Datastore Architecture

ASCII Art mostly from RFC 8342
 Each datastore's contents defined by YANG.



- If the device has no configuration and only has operational state that can be monitored, then it only needs to present the “operational” datastore.
- If the device has a configuration, then it becomes necessary for it to have the “running” configuration datastore. There are different types of configuration datastores: the “factory-default” datastore, which represents the configuration the device has when it is shipped from the manufacturer’s factory. There’s also a “startup” datastore, which is a user-specified or configurable datastore representing what is the configuration for the device when it reboots. There’s also a “candidate” datastore; this datastore is sort of a sandbox datastore where configuration updates can be accumulated over time, validated, and ultimately merged into the “running” datastore whenever the running datastore is updated.
- The intended datastore reflects the desired configuration state, and the operational datastore contains the actual state of the device as it operates. YANG validation, covered later in this paper, occurs on the intended datastore. Applying the configuration dictated by this datastore is subject to local factors, such as what hardware is present on the device at the time and other delays that may occur.
- The operational datastore is what a device’s current configuration is. This may not be what was configured, but the operational datastore is what’s actually implemented by the data plane.

NDMA provides a structured approach to handling network configuration and state data. NMDA is designed to work seamlessly with network management protocols like NETCONF and RESTCONF. These protocols can interact with the various datastores defined by NMDA, enabling fine-grained control over network configurations and state data, and supporting model-driven network management.

NETCONF and RESTCONF Protocols

NETCONF and RESTCONF are two key protocols that leverage NMDA to provide robust network management capabilities. NETCONF, exclusively XML-based, offers extensive features for managing configurations, including capabilities for transactional operations, rollbacks, and validation. RESTCONF, on the other hand, provides an HTTP-based interface that supports both XML and JSON encodings, making it more accessible for modern application development. Despite these differences, both protocols are designed to work seamlessly with YANG data models, ensuring that configuration and state data are managed consistently and accurately across the network.

	NETCONF	RESTCONF
Encodings	XML-only	XML, JSON, ++
Transports	SSH or TLS	HTTP, over TLS or QUIC
Operations	RPC-based (<edit-config>, etc.)	REST-based (GET, PUT, PATCH, etc.)
YANG	Fully Aware	Fully Aware
NMDA	Fully Supports	Fully Supports
Transactional	Yes	Yes
Status	Entrenched	Easier (curl, postman, etc.)

Transactional API

- Any valid configuration can move to any other valid configuration in a single protocol operation

When comparing NETCONF and RESTCONF, several key similarities and differences emerge. Both protocols support operations targeting individual nodes within the YANG data tree, allowing precise and granular management of network configurations. They are fully YANG-aware, meaning they can interpret and manipulate data defined by YANG models. However, NETCONF operates over SSH and TLS, providing a secure channel for configuration changes, and uses custom XML-based operations. RESTCONF, in contrast, leverages standard RESTful operations (GET, POST, PUT, PATCH, DELETE) over HTTP, which are more intuitive for developers familiar with web technologies. This makes RESTCONF easier to implement and integrate, especially in environments where JSON is preferred for its lightweight and readable format.

The primary differences between NETCONF and RESTCONF lie in their transport and operational mechanisms. NETCONF's use of SSH and TLS ensures secure communication, which is crucial for maintaining the integrity of configuration changes. Its custom XML-based operations, such as <edit-config> and <commit>, provide robust transactional capabilities, ensuring that configuration changes can be validated and rolled back if necessary.

RESTCONF, on the other hand, offers a more straightforward and flexible interface, making it easier to integrate with modern web applications and tools. Its reliance on standard HTTP methods and support for both XML and JSON formats enhance its accessibility and usability.

YANG Data Modeling

YANG, the data modeling language used by both NETCONF and RESTCONF, plays a crucial role in defining the structure and semantics of configuration and state data. YANG allows network operators to create comprehensive data models that describe the

configuration and operational aspects of network devices. One of YANG's strengths is its ability to clearly distinguish between configuration data (the intended state of the network) and operational state data (the actual state of the network). This distinction is essential for effective network management, as it enables operators to understand not only what the desired configurations are but also how they are currently being implemented on the devices.


YANG (Yet Another Next Generation)

Aliro™

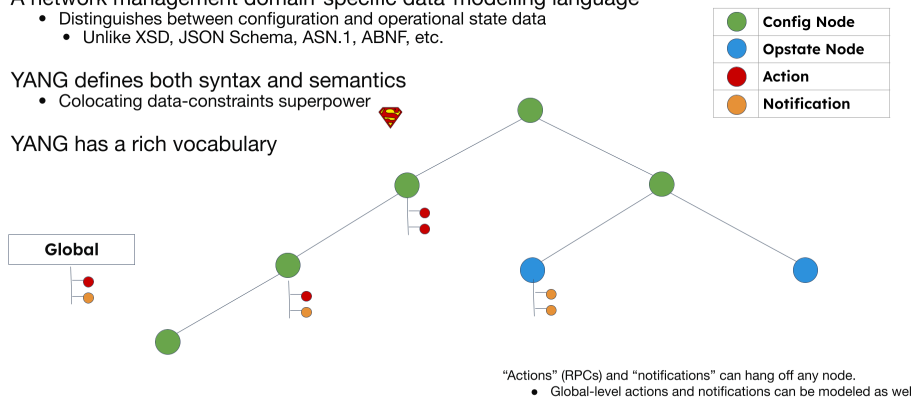
A network management domain-specific data-modelling language

- Distinguishes between configuration and operational state data
- Unlike XSD, JSON Schema, ASN.1, ABNF, etc.

YANG defines both syntax and semantics

- Colocating data-constraints superpower 

YANG has a rich vocabulary



In YANG, configuration nodes can only have other configuration nodes as their ancestors, ensuring a clear hierarchy and relationship between different pieces of configuration data. This structured approach helps prevent misconfigurations and ensures that changes are applied systematically. Operational state nodes, on the other hand, can have either configuration or operational state nodes as their ancestors. Operational state includes, e.g., real-time status information and statistics about the network's performance. By co-locating data constraints with data definitions, YANG provides a powerful framework for ensuring data integrity and consistency, making it an indispensable tool for modern network management.

TMN Layering Model and Management Systems

Introduction to TMN Layering

The Telecommunications Management Network (TMN) layering model is a structured approach to managing telecommunications networks. Developed by the International Telecommunication Union (ITU), the TMN model aims to standardize network management practices, with a clear hierarchy of management functions and layers. This model enhances interoperability, scalability, and efficiency by clearly delineating the roles and responsibilities of different management layers, ensuring that network operations are conducted smoothly and systematically.

Different Layers and Their Functions

The TMN layering model comprises four primary layers: the Element Management System (EMS), the Network Management System (NMS), the Service Management System (SMS), and the Business Management System (BMS). Each layer has distinct functions and responsibilities, working together to provide comprehensive management of the network. The EMS is responsible for managing individual network elements, such as routers and switches, ensuring they operate correctly and efficiently. The NMS oversees the network as a whole, presenting a normalized data-model abstracting differences between specific devices. The SMS bridges the gap between user services and the network infrastructure, managing services like VPNs, QoS, and other value-added services. The BMS focuses on the business aspects, including billing, customer management, and service level agreements (SLAs). The graphic below shows these layers in a pyramid. This is intended to convey that there are many more Network Elements than there are Element Management Systems. On the lefthand side are protocols commonly used for that layer. NETCONF is the protocol commonly used for the Network Elements to present to the management system. The various management system layers, because they don't have to use NETCONF as much as the network elements do, can use REST-based APIs. Any reasonable programmatic API would be fine, but REST-based APIs are very commonly used today for web-applications. When using a REST API and leveraging YANG in the management system, it makes sense to use RESTCONF as well, especially if the goal is to convey that YANG data up to higher levels of the management system stack.

Detailed breakdown of each layer for classical and entanglement-based quantum networks


At the base of the TMN model is the Network Element, which handles the “data plane” traffic of the network. In classical networks, the data plane would be IP. In entanglement-based Quantum Networks the data plane is light, free space links, or optical fiber.

The Element Management System (EMS) handles fault management, configuration, accounting, performance, and security (FCAPS) at the device level. It serves as the interface between the physical network components and higher-level management systems, ensuring that each element operates optimally and any issues are addressed promptly. Typically, the EMS manages the resources or the kinds of devices that were produced by a single vendor. It has few to no abstractions over the network element-provided data models. The EMS collects data from network elements, performs diagnostics, and implements configuration changes. Many times the element management system is hierarchically and geographically distributed, because it's necessary to deploy data collection nodes where the devices are located in order to collect data efficiently without consuming too much network

bandwidth. In entanglement-based networks, the EMS would manage quantum-specific devices like quantum routers, quantum repeaters, and entangled photon sources.

TMN “Layering” Model

Aliro™



- **BMS: Business Management System**
 - System managing functions related to a specific business
 - Often developed in-house
- **SMS: Service Management System**
 - System managing “services” in the network
 - Uses high-level abstractions to bridge gap between what users want and the network implements. E.g., QoS, Firewall, VPN, attack detection/mitigation/prevention, compute, and network
- **NMS: Network Management System**
 - System managing a network of different kinds of resources
 - Uses abstractions to generalize EMS-specific data-models
- **EMS: Element Management System**
 - System managing a single kind of resource
 - Few to no abstractions over NE provided models
 - Many times hierarchically and geographically distributed
 - Typically vendor-specific (supplied by vendor)
- **NE: Network Element**
 - Devices handling the “data plane” traffic of the network

The Network Management System (NMS) operates at a higher level, overseeing the network’s overall performance and health. It manages the interconnections between network elements, ensuring seamless data flow and communication. The NMS is responsible for network-wide fault detection and resolution, performance monitoring, and capacity planning. It provides a holistic view of the network, allowing administrators to make informed decisions about network optimization and expansion. In an entanglement-based Quantum Network, the NMS would manage entanglement distribution across the network, ensuring that entanglement links maintain high fidelity and low error rates. It would handle the orchestration of entanglement swapping operations and the routing of quantum information through the network, optimizing the usage of entanglement-based resources.

The Service Management System (SMS) translates user requirements into network services. The SMS ensures that the network can bridge the gap between what users want and what the network implements. For instance, a network may have services such as quality of service (QoS), firewall, VPN, attack detection/mitigation/prevention, compute-as-a-service, network-as-a-service, etc. There are many different kinds of services that could be present, and even domain-specific. In the context of entanglement-based Quantum Networks, the SMS would manage quantum-specific services such as key generation as a service, entanglement as a service, and teleportation as a service.

At the top of the TMN hierarchy, the Business Management System (BMS) is managing the functions related to a specific business, and often it’s developed in-house because it’s very business-specific. However, there are some tools that businesses can purchase and customize or have customized for them in order to implement their business management system. The BMS performs functions related to the network as a business by analyzing

quality issues, and even providing a basis for billing and other financial reports. This is true for both classical networks and entanglement-based Quantum Networks.

By clearly defining the roles and responsibilities of different management layers, the TMN model enhances the efficiency, scalability, and reliability of network operations.

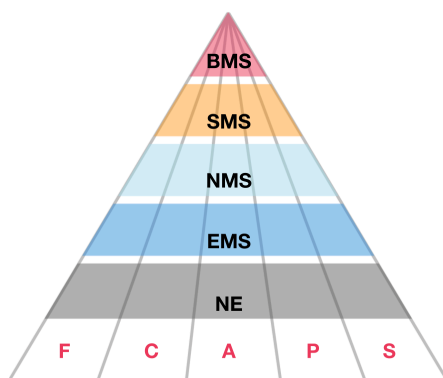
FCAPS Model

The FCAPS model is used within the TMN layering model, providing a framework for managing network operations at different layers of the TMN model. FCAPS stands for Fault, Configuration, Accounting, Performance, and Security management, each of which is essential for maintaining a network.

- Fault management involves detecting, logging, and resolving network issues.
- Configuration management gathers, stores, and tracks configurations from the network devices. It also updates and pushes configurations to the networking devices, and many times it includes inventory management, software management, and license management as well.
- Accounting management tracks network usage for billing and auditing purposes, and for enforcing quotas.
- Performance management ensures that the network meets performance standards and optimizes resource utilization. It supports capacity planning and service level agreements.
- Security management controls access to network resources according to local policy. It protects the network from unauthorized access and ensures data integrity. The Triple A (Authentication, Authorization, and Accounting/Auditing) is performed at this layer. Certifications such as FIPS or Common Criteria would be performed on the security management layer.

TMN “FCAPS” Model

Aliro™



Kinds of Management

Fault Management

- Recognize, log, and correct (to the extent possible) faults occurring in the network.

Configuration Management

- Gather, store, track configurations from network devices.
- Update/push configurations to network devices.
- Includes Inventory, Software and License management as well.

Accounting Management

- Track network utilization information for billing purposes
- Also used to enforce quotas.

Performance Management

- Ensure network performance is maintained at acceptable levels.
 - E.g., throughput, response times, utilization, etc.
- Supports capacity planning and service-level agreements (SLAs)

Security Management

- Control access to network resources according to local policy.
- Authentication, Authorization, and Accounting/Auditing (AAA)
- Certifications: FIPS and Common Criteria

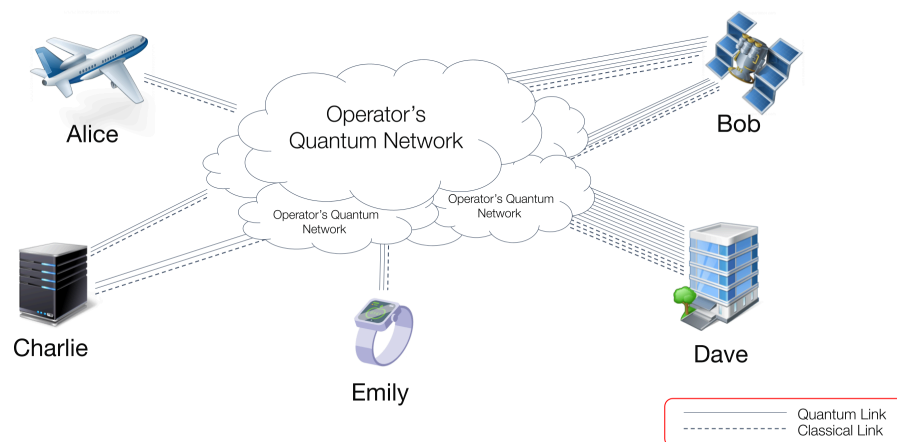
The FCAPS model ensures that all critical aspects of network management are addressed systematically through different layers of the network.

TMN Example Topology

In our network topology example, we have several entanglement-based nodes labeled Alice, Bob, Charlie, and Dave. These labels do not denote individuals but rather pieces of sophisticated network equipment, as shown in the provided slides. Each entanglement-based node is interconnected within the operator networks using both quantum and classical links. The diagram illustrates that some nodes possess multiple quantum-classical links, which are essential for maintaining network robustness and ensuring high availability through redundancy. The diagram also depicts a satellite with links connecting to two different operator networks. This configuration allows the satellite to function as a repeater, extending the network's reach and facilitating communication between disparate parts of the network. Although the diagram portrays these nodes as external to the operator network for clarity, they are, in practice, integral components of the network infrastructure. This setup highlights the versatility and scalability of modern network architectures, where seamless integration across various segments is crucial.

TMN Example: Topology

Aliro™



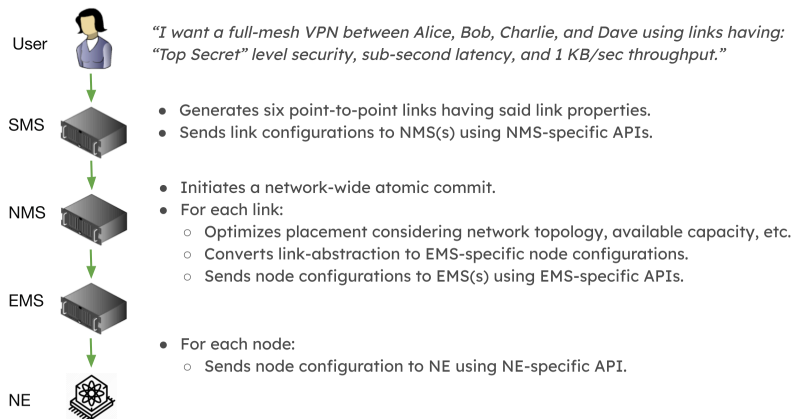
The current scale of entanglement-based Quantum Network equipment is substantial, akin to the early days of mainframe computers that occupied entire rooms or buildings. Each entanglement-based node requires an array of devices, such as optical tables, measure modules, detectors, timing modules, lasers, and polarization correctors. This extensive setup underscores the complexity and space requirements of today's entanglement-based networks. However, there is a promising vision for the future where this equipment could be miniaturized to the size of a smartwatch, represented by Emily in the image, significantly reducing the physical footprint and enabling more flexible deployment options. What's truly interesting about this possibility is it could help solve the identity problem: how do we know

that Alice is really Alice or Charlie is really Charlie? Well, they're actually wearing a device on their personal self. Then there may be some ability for using quantum identity to know for sure that you're speaking with a certain person.

TMN Example: Configuration

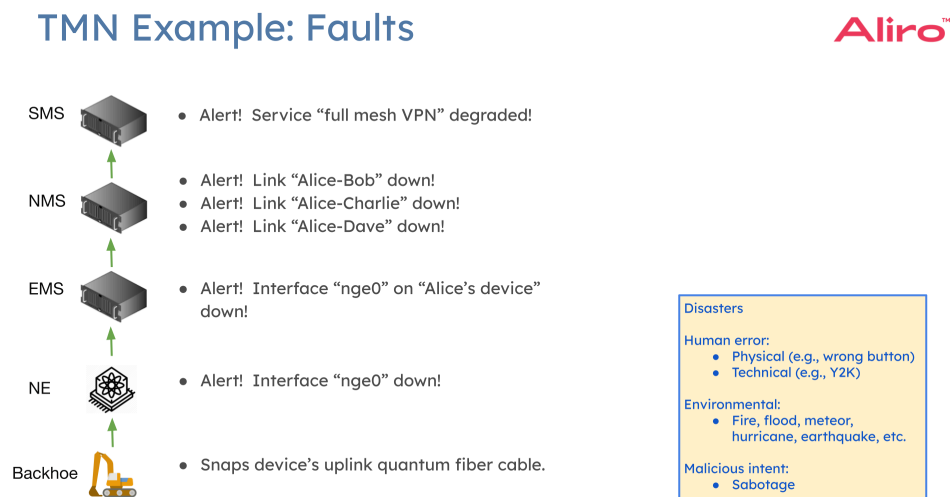
When it comes to configuration management, consider a user requesting a full-mesh VPN between Alice, Bob, Charlie, and Dave with specific requirements for top-secret security, sub-second latency, and 1 KB/sec throughput. As illustrated in the slides, the Service Management System (SMS) processes this request by generating the required point-to-point links with the specified properties. These configurations are then sent to the Network Management Systems (NMS) through NMS-specific APIs. The NMS is responsible for initiating a network-wide commit, optimizing link placement, and converting link abstractions into Element Management System (EMS) specific configurations, ensuring the network meets the user's demands. In this example, we see how the "configuration" attribute (the 'C' in FCAPS) is going through the different layers, and it has different meaning in each of those layers.

TMN Example: Configuration



TMN Example: Faults

There are many reasons a fault may occur. There's human error, for instance, pressing the wrong button; technical errors such as Y2K; environmental issues such as fire, flood, meteor, hurricane, earthquake. Faults can even be caused with malicious intent, like sabotage. In fault management, consider a scenario where a physical disruption, such as a backhoe snapping an optical fiber cable used to carry qubits, occurs. The affected network element detects this event and sends an alert indicating the interface is down, which is received by the EMS. The EMS escalates this to a higher-level alert, specifying that Alice's device interface is down. The NMS, realizing the broader network impact of this alert, generates additional alerts for each disrupted link connected to Alice's device, as depicted in the slides. This multi-layered alert system ensures that faults are quickly identified and addressed, maintaining the network's operational integrity.



Self-Healing Networks

What can be done in scenarios such as the fault example above? The concept of self-healing networks is crucial for maintaining service levels and adhering to SLAs. When a disruption is detected, such as the downed link in the previous example, the network's self-healing mechanisms kick in. These mechanisms automatically identify the issue, determine the best corrective action, and implement changes to restore service. For instance, the system might re-route traffic or reconfigure network elements to bypass the fault. This approach minimizes downtime and ensures continuous service delivery. This is critical because SLAs often require high levels of network availability and performance, commonly referred to as "five nines" availability, which allows for only a few minutes of unscheduled downtime per year.

Achieving this level of reliability necessitates robust network architectures, redundancy, and advanced fault management strategies. The integration of quantum technologies and

self-healing capabilities, as described in the slides, plays a vital role in meeting these stringent requirements. By ensuring that networks can quickly recover from faults and maintain high performance, we can deliver the reliable services expected by users and operators.

Managing Entanglement-based Quantum Networks

Managing entanglement-based Quantum Networks involves addressing a range of unique challenges that stem from the fundamental differences between quantum and classical communication. Despite these differences, the principles of effective network management remain consistent, focusing on reliability, security, and performance. As the field of entanglement-based networking continues to develop, ongoing research and standardization efforts will be crucial in overcoming these challenges and achieving efficient and scalable entanglement-based Quantum Network management.

Quantum network management vs. classical network management

By and large, managing entanglement-based Quantum Networks is similar to managing classical networks. The TMN Layering and FCAPS models still hold, the nonfunctional attributes remain unchanged, and functional attributes are similar enough that they only need slight shifting to be relevant in the entanglement-based communication domain. However, when comparing entanglement-based and classical network management practices, several key differences emerge:

No physical topology discovery. In the classical networking world, topology is discovered. Topology discovery protocols have been developed to determine how ports are connected. LLDP, or the Link Layer Discovery Protocol, is the most widely used. By topology, this regards how the various physical devices are connected to each other, e.g., via their ports. If you have a device with four ports and different ports have fibers connecting to different other equipment, the question is, what other port on the other equipment is this port connected to? One of the primary challenges in entanglement-based network orchestration is topology discovery.

Entanglement-based networks lack such standardized protocols as LLDP, making it difficult to determine how different entanglement-based nodes are connected. This challenge is exacerbated by the fact that quantum communication fibers (optical fiber that carries qubits) are typically fully utilized to convey qubits, leaving no room for the transmission of metadata required for topology discovery. In classical networking, the ubiquitous presence of wave division multiplexing was used to convey metadata. So already there is wave division multiplexing, and adding another virtual channel for conveying topology was easy to do. However, in the entanglement-based networking world, fibers are generally fully utilized to convey the data or the qubits, with no provision to convey the metadata. If it's not possible to convey the metadata in-band,

then an overlay classical network must be utilized to convey that metadata. In either case, agreement is needed among the quantum device vendors to achieve some form of physical topology discovery for optical networks.

Need for protocol adaptors. In the classical networking devices, there was a long period of needing protocol adapters, as each vendor had proprietary CLIs and a number of scripts were written to interoperate with the CLIs. SNMP didn't get widely used, but NETCONF is firmly established in the classical networking industry, such that protocol adapters are now for the most part a thing of the past in the classical networking world. However, quantum devices being at an early stage, and standardization is not yet in place. The lack of standardization means that each device may require a custom protocol adapter to interface with the network management system. As the industry matures, there is a strong need for using standardized protocols (e.g., NETCONF) to facilitate interoperability and simplify network management. Until then, the reliance on protocol adapters remains a necessary but cumbersome aspect of managing entanglement-based networks.

Need for stabilization delays. Next difference is the need for stabilization delays. Some quantum devices require significant time to apply configuration changes. For instance, a laser may require up to 5 minutes to reach temperature. Needing to wait for the intended configuration to be applied is not new, but the wait time is an order of magnitude more. Five minutes is a long wait. Entanglement-based orchestration systems managing such devices need to account for stabilization when provisioning services in a timely manner.

Need to route timing signals along the data path. There's also another difference - the need to route timing signals the same as the photons, as quantum communication fibers (optical fiber that carries qubits) typically only convey qubits with no wave division multiplexing. The timing protocol signals must be routed over alternate cables, which may be either fiber or copper. However, some timing protocols rely on measuring the light transmission propagation delay and thus must be routed over similar fibers along the same path in order to mimic the qubit delay within sub-nanosecond accuracy. Entanglement-based orchestration systems managing timing models using such protocols must ensure that both signals are routed together.

Need to periodically calibrate polarization. Many quantum devices are polarization sensitive. They assume horizontal, vertical, diagonal, and anti-diagonal polarizations. Yet polarization can drift throughout the day. Environmental factors like temperature change can cause polarization to drift. Polarization-maintaining fiber helps, but is expensive, and thus only used for short runs (e.g., inside a data center). Polarization compensators can be inserted into the data path, but they need to be

adjusted periodically, for instance, once per hour. Entanglement-based orchestration systems managing networks containing APCs, or automatic polarization compensating devices, need to ensure that they are calibrated periodically. This brings up an interesting point we touched on earlier about the utilization of the network. While certain services are up and running, there may be some capacity that's not being utilized. While it's not being utilized, the system could be calibrating the polarizers so that they're ready when needed.

Despite these challenges, the principles of Quantum Network management share many similarities with classical network management, particularly in the overarching goals of ensuring reliability, security, and performance. Both types of networks require comprehensive management of configuration, fault detection, performance monitoring, and security. However, the methods and technologies employed to achieve these goals differ significantly, reflecting the unique properties and requirements of entanglement-based networks. Entanglement-based networks require software and hardware components that mirror those of classical networks, but that operate using qubits and entanglement rather than bits.

Conclusion

Entanglement-based quantum networks are being built today by a variety of organizations for a variety of use cases – benefiting organizations internally, as well as providing great value to an organization's customers. Telecommunications companies, national research labs, intelligence organizations, and systems integrators are just a few examples of the organizations Aliro is helping to leverage the capabilities of Quantum Networking.

Building quantum networks that use entanglement is no easy task. It requires:

- Emerging hardware components necessary to build the network.
- The software necessary to design, simulate, run, and manage the network.
- A team with expertise in the fundamental science of entanglement-based quantum networks and classical networking.
- Years of hard work and development.

This may seem overwhelming, but Aliro is uniquely positioned to help you build your Quantum Network. The steps you can take to ensure your organization is meeting the challenges and leveraging the benefits of the secure networking revolution are part of a clear, unified solution already at work in networks like the EPB Quantum NetworkSM powered by Qubitekk in Chattanooga, Tennessee.

AliroNetTM, the world's first full-stack entanglement-based network solution, consists of the software and services necessary to ensure customers will fully meet their quantum secure

networking goals. Each component within AliroNet™ is built from the ground up to be compatible and optimal with entanglement-based networks of any scale and architecture. AliroNet™ is used to simulate, design, run, and manage Quantum Networks as well as test, verify, and optimize entanglement-based hardware for network performance. AliroNet™ leverages the expertise of Aliro personnel in order to ensure that customers get the most value out of the software and their investment.

Depending on where customers are in their Quantum Networking journeys, AliroNet™ is available in three modes that create a clear path toward building full-scale entanglement-based secure networks: (1) Emulation Mode, for emulating, designing, and validating entanglement-based networks, (2) Pilot Mode for implementing a small-scale entanglement-based network testbed, and (3) Deployment Mode for scaling entanglement-based networks and integrating end-to-end applications. AliroNet™ has been developed by a team of world-class experts.

To get started on your Quantum Networking journey, reach out to the Aliro team for additional information on how AliroNet™ can enable secure communications.

www.alirotech.com

References

"FCAPS." Wikipedia, <https://en.wikipedia.org/wiki/FCAPS>.

"IETF | Internet Engineering Task Force." IETF, <https://www.ietf.org>.

"NETCONF." Wikipedia, <https://en.wikipedia.org/wiki/NETCONF>.

"REST." Wikipedia, <https://en.wikipedia.org/wiki/REST>.

"Simple Network Management Protocol." Wikipedia, https://en.wikipedia.org/wiki/Simple_Network_Management_Protocol.

"Telecommunications Management Network." Wikipedia, https://en.wikipedia.org/wiki/Telecommunications_Management_Network.

"YANG." Wikipedia, <https://en.wikipedia.org/wiki/YANG>.

"RFC 6241 - Network Configuration Protocol (NETCONF)." IETF, <https://datatracker.ietf.org/doc/html/rfc6241>.

"RFC 8040 - RESTCONF Protocol." IETF, <https://datatracker.ietf.org/doc/html/rfc8040>.

"RFC 8342 - Network Management Datastore Architecture (NMDA)." IETF, <https://datatracker.ietf.org/doc/html/rfc8342>.