Multi-Network Enterprise Solutions

ATIS-I-0000086
July 2021

atis

COPYRIGHT INFORMATION

ATIS-I-00000XX

Alliance for Telecommunications Industry Solutions

1200 G Street, NW, Suite 500

Washington, DC 20005

NOTICE OF DISCLAIMER AND LIMITATION OF LIABILITY

The information provided in this document is directed solely to professionals who have the appropriate degree of experience to understand and interpret its contents in accordance with generally accepted engineering or other professional standards and applicable regulations. No recommendation as to products or vendors is made or should be implied.

NO REPRESENTATION OR WARRANTY IS MADE THAT THE INFORMATION IS TECHNICALLY ACCURATE OR SUFFICIENT OR CONFORMS TO ANY STATUTE, GOVERNMENTAL RULE OR REGULATION, AND FURTHER, NO REPRESENTATION OR WARRANTY IS MADE OFMERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR AGAINST INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. ATIS SHALL NOT BE LIABLE, BEYOND THE AMOUNT OF ANY SUM RECEIVED IN PAYMENT BY ATIS FOR THIS DOCUMENT, AND IN NO EVENT SHALL ATIS BE LIABLE FOR LOST PROFITS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES. ATIS EXPRESSLY ADVISES THAT ANY AND ALL USE OF OR RELIANCE UPON THE INFORMATION PROVIDED IN THIS DOCUMENT IS AT THE RISK OF THE USER.

NOTE - The user's attention is called to the possibility that compliance with this document may require use of an invention covered by patent rights. By publication of this document, no position is taken with respect to whether use of an invention covered by patent rights will be required, and if any such use is required no position is taken regarding the validity of this claim or any patent rights in connection therewith. Please refer to [http://www.atis.org/legal/patentinfo.asp] to determine if any statement has been filed by a patent holder indicating a willingness to grant a license either without compensation or on reasonable and non-discriminatory terms and conditions to applicants desiring to obtain a license.

ABSTRACT

Enterprises have the freedom to use a wide variety of networks and network technologies, including public and private cellular, office WLANs and public Wi-Fi hotspot services, SD-WANs and application-specific networks such as LoRa for IoT. But this freedom comes with a host of challenges. For example, each network often has different authentication credentials with separate identity providers, as well as different security capabilities, policies, levels of performance and coverage domains. All of these variables create complexity, which drives up the cost of managing them.

ATIS recognized these challenges and identified alternatives to integrate and simplify the management of this complex network environment, as well as enhance the overall reliability, resiliency, performance and security of the end-to-end system. This report describes those challenges and the technologies that enterprises can use to mitigate them. One example is the use of multipath to better leverage multiple networks in concert to increase overall network reliability, resiliency and performance.

FOREWORD

As a leading technology and solutions development organization, the Alliance for Telecommunications Industry Solutions (ATIS) brings together the top global ICT companies to advance the industry's business priorities. ATIS' 150 member companies are currently working to address network reliability, 5G, robocall mitigation, smart cities, artificial intelligence-enabled networks, distributed ledger/blockchain technology, cybersecurity, IoT, emergency services, quality of service, billing support, operations and much more. These priorities follow a fast-track development lifecycle from design and innovation through standards, specifications, requirements, business use cases, software toolkits, open source solutions and interoperability testing.

ATIS is accredited by the American National Standards Institute (ANSI). ATIS is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a founding Partner of the oneM2M global initiative, a member of the International Telecommunication Union (ITU), as well as a member of the Inter-American Telecommunication Commission (CITEL). For more information, visit www.atis.org. Follow ATIS on Twitter and on LinkedIn.

COPYRIGHT INFORMATION

ATIS-I-0000086

Alliance for Telecommunications Industry Solutions

1200 G Street, NW, Suite 500

Washington, DC 20005

NOTICE OF DISCLAIMER AND LIMITATION OF LIABILITY

TABLE OF CONTENTS

1.
EXECUTIVE SUMMARY

Enterprises use a wide variety of networks and network technologies to deliver services and support internal operations. Each network may have different authentication credentials with separate identity providers, as well as different security capabilities, policies, levels of performance and coverage domains. All of this creates a network environment that is complex to operate and manage. As a result, network performance and efficiency may not be optimized, leading to decreased productivity for enterprise users.

This report explores an industry approach designed to integrate and simplify the management of this complex network environment and enhance the overall reliability, resiliency, performance, and security of the end-to-end (E2E) system. To better illustrate these challenges, the report uses various user stories to highlight and identify specific gaps that need to be addressed, including:

- Identity management to simplify authentication for both network and application access.
- The use of multipath technologies to better leverage multiple networks in concert to increase overall network reliability, resiliency and performance. These technologies often require a robust policy solution to propagate policies across networks and applications.
- Enterprise security aspects such as timing resiliency via Position Navigation and Timing (PNT) and Secure Access Service Edge (SASE [1]) cloud-delivered security services.

Deploying a standardized platform enables enterprises to more efficiently manage their various networks, which benefits both the enterprise and network providers. SD-WAN and cloud-based solutions are used as example platforms. By deploying new, multi-network-enabling technologies on existing platforms, enterprises reduce management costs while increasing security and performance. In addition, network operators can offer more advanced services and capabilities that interface directly with the platform and can potentially offer the platform itself as a service to the enterprise.

1 SASE is a cloud-delivered service model that identifies users and devices to apply policy-based security delivering secure access to the appropriate application or data.

2. INTRODUCTION

This report defines “enterprise” as an entity that:

- Has a set of consumers, users, employees and/or members.
- Delivers a service or application to those consumers, members, etc.
- Uses one or more communications paths, each of which might comprise multiple communications providers.

Examples of enterprises as used here could include carpeted enterprises, private industrial networks and entities focused on delivering an application, public or private, as all or part of their business. These could include large banks or insurance companies, hyperscale providers of goods and services such as Amazon, large-scale providers of transportation and delivery services such as UPS or providers of communications and collaboration services such as Zoom.

Enterprises use a wide variety of networks and network technologies to deliver services and support internal operations, such as:

- Mobile wireless services from multiple mobile operators based on both employee and end-customer preferences using multiple 3GPP technologies (3G, 4G and 5G).
- Private (enterprise) Wi-Fi in dedicated enterprise spaces.
- Private 3GPP networks (e.g., LTE, 5G with CBRS and other configurations).
- Public Wi-Fi.
- Wide-area low-power (IoT-specific) technologies (e.g., LoRa).
- Wired/fixed access and transport services (e.g., in support of enterprise VPNs and dedicated data center interconnect via fiber, Ethernet services, DOCSIS).
- Satellite-based communication services.
- Public cloud and/or SASE-delivered networking/security service to connect large offices and branch offices.
- Other new/developing communication services and technologies, including the need for accurate and redundant timing sources.

Operating and managing this diverse communication environment is challenging. Each network may have different authentication credentials with separate identity providers, as well as different security capabilities, policies, levels of performance and coverage domains. Enterprises often depend on the “least common denominator” for network capabilities while building complex, bespoke overlay solutions to achieve an acceptable level of management control. Inevitably, the result is an expensive and complex environment with sub-optimal E2E network performance. In addition, this least-common-denominator approach limits the types of features and services offered by communication providers.

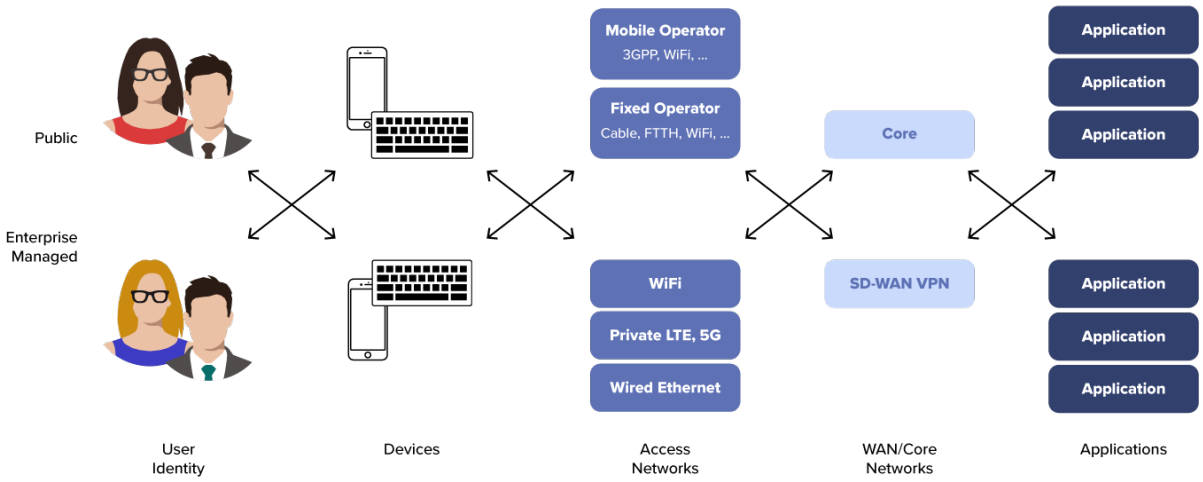

Figure 2.1 – Multi-Network Enterprise Complexity

Enterprise networking must handle a wide variety of multi-network options, as illustrated in Figure 2.1. For example:

- An enterprise user could be an employee, a contractor with enterprise credentials or a public/non-enterprise user.
- The user may need to access enterprise services via their personal device, a public device or an enterprise-managed device.
- This device may attach to a public mobile or fixed access network or to an enterprise managed network (wireless or fixed).
- Network access to applications may be via a public core network or via an enterprise-managed VPN technology (or SD-WAN).
- Applications may exist in the public domain or enterprise domain and may be deployed using a variety of cloud models (e.g., SaaS, PaaS, IaaS).

This multiplicity of options makes it challenging for the enterprise to provide needed services to users with the right performance and security in a cost-effective and manageable way.

Addressing these challenges is an industry priority, given the existing market size and high growth potential. In addition to the existing enterprise market needs related to service delivery and support for internal operations, the vertical market segment is expected to expand significantly.

This report explores an industry approach designed to integrate and simplify the management of this complex network environment and enhance the overall reliability, resiliency, performance and security of the E2E system.

Specifically, this work will:

- Articulate various user stories to highlight and identify specific gaps that need to be addressed.
- Identify the enabling technologies that can be leveraged to address the problem space. These include areas such as:
  - Mechanisms to distribute a common, enterprise-controlled policy set to all network domains in a consistent manner.
  - Common authentication and identity management structures to simplify authentication across all of the various network modes.
  - Security capabilities to unify and enhance overall security across multiple networks and types.

    - Network capabilities such as multipath to better leverage multiple networks in concert to increase overall network reliability, resiliency and performance.
    - Support for transparent onboarding across multiple networks and technologies.
- Investigate needed solutions in the space including leveraging enabling technologies to support and supplement the identified gaps.

Deploying a standardized platform enables the enterprise to more efficiently manage its various networks, offering benefits to both the enterprise and network providers. In this way,

the enterprise achieves lower overall management costs while increasing security and performance. In addition, network operators can offer more advanced services and capabilities that interface directly with the platform and can potentially offer the platform itself as a service to the enterprise.

Realizing this vision requires the industry — vendors, operators, and enterprises — to work collaboratively to develop a consensus view of the requirements and framework, identify and document the elements missing from communication technology standards and drive adoption of the concepts necessary for integrated connectivity.

3
USER STORIES

3.1
MULTI-NETWORK POSITION, NAVIGATION AND TIMING (PNT) RESILIENCY

GPS meets many of today's enterprise PNT needs. Precision timing signals sent through GPS are used to synchronize cellphone calls, time-stamp financial transactions, support safe travel by aircraft, ship, train and car and enable other dedicated industrial communication functions. However, GPS transmissions can be disrupted unintentionally by radio interference or the weather in space and intentionally through jamming or spoofing. As such, there is increased interest in enhancing its resiliance.

On February 12, 2020, US Executive Order 13905 was released on "Strengthening National Resilience through Responsible Use of Positioning, Navigation, and Timing Services." [2]

According to the order: "The widespread adoption of PNT services means disruption or manipulation of these services could adversely affect U.S. national and economic security. To strengthen national resilience, the Federal Government must foster the responsible use of PNT services by critical infrastructure owners and operators."

In February 2021, the National Institute of Standards and Technology (NIST) published "NISTIR 8323 - Cybersecurity Profile for the Responsible Use of Positioning, Navigation, and Timing (PNT) Services" to help organizations make deliberate, risk-informed decisions about their use of PNT services.[3]

In addition, the ATIS SYNC committee has published a number of documents related to network timing, including the ATIS SYNC Technical Report (ATIS-0900005) on GPS Vulnerability, which describes the telecommunications industry's dependence on GPS and highlights GPS vulnerabilities of concern to the communications sector. [4] ATIS SYNC recommends that telecom carriers explore options including time-sync networks engineered to provide time as a service both internally and externally.

- Relative to enterprise positioning services, the FCC has created requirements to address location accuracy for emergency calls. These rules are discussed in the FCC 4th Report & Order and codified within the published Code of Federal Regulations (CFR). ATIS standard ATIS-0700028.v002, *Location Accuracy Improvements for Emergency Calls (version 2)* describes the standards needed to support the commitments defined in the roadmap described above, as well as the rules as outlined within the FCC CFR.

3.1.1 Story Highlights

An enterprise has the need for precision timing to support ongoing business operations. This need can include:

- Operation of private LTE/5G networks that may need precision timing.
- Time stamping of critical transactions such as financial transactions.
- Other industrial timing applications and needs.

2 https://www.federalregister.gov/documents/2020/02/18/2020-03337/strengthening-national-resilience-through-responsible-use-of-positioning-navigation-and-timing

3 https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8323.pdf

4 https://access.atis.org/apps/group_public/download.php/36304/ATIS-0900005.pdf

GPS often is leveraged as a timing source. However, GPS is particularly susceptable to jamming and spoofing attacks. GPS jammers are typically small, self-contained transmitters. Although illegal, they are widely available for purchase online and often cost as little as $100. Their relatively low power and quick start-up time enable them to be used only when required. When active, the GPS jammer generates an interference signal over a 5 to 10 meter radius to disrupt reception of the GPS signal and thus disrupt the precision timing source.

In addition, satellite malfunction, atmospheric disturbances or solar flares can temporarily disrupt the transmission of GPS signals. When this occurs, the GPS device used for precision timing may no longer be able to acquire or maintain timing synchronisation.

These drawbacks and vulnerabilities highlight the need to provide more resilient sources for precision timing in the enterprise.

3.1.2 Business Drivers

For businesses with a critical business need for precision timing, the ease with which GPS-based timing can be disrupted means that a resilient precision timing architecture is essential for miminizing timing-related business disruption and associated cost and revenue implications.

3.1.3 Actors

The primary actor is the enterprise itself. Each enterprise may or may not have critical business functions dependent on precision timing, so an internal investigation is warranted to determine whether there is or is not a need for precision timing. If so, depending on the chosen solution for resiliency, many different vendors and service providers may be actors. Within the scope and context of this report, key actors would be network service providers that offer precision timing services via their interconnect services.

3.1.4 Deployment Model and High-Level Architectural Context

Precision timing services can be provided on top of many forms of connectivity. Section 7 of ATIS-0900005: ATIS SYNC Technical Report on GPS Vulnerability outlines a variety of methods to enhance precision timing resiliency. Trials transporting very high precision time and phase synchronization over fiber using IEEE-1588v2 PTP have shown promise. PTP packetizes time and phase information for delivery over a packet-based network such as Ethernet, which is in turn transported over fiber. In addition, ITU-T standard J.211 describes a two-way protocol transported over the physical layer that includes a mechanism to correct for

transport delay and asymmetry. It is not packet based and thus is not impaired by delay variation. As such, precsion timing as a service seems feasible. This enables highly resilient precision timing solutions by leveraging one or more network service provider precision timing services in a multi-network enterprise solution.

3.2 FLEXIBLE MULTI-NETWORK, MULTI-APP IDENTITY FEDERATION WITH VERIFIABLE CHARGING MECHANISMS

Many enterprises are challenged by the need to support employee authentication credentials across a wide variety of networks and applications. Use of identity-federation mechanisms can provide some relief, but many such mechanisms do not support the wide variety of systems that require authentication.

For example, the enterprise must address two classes of authentication:

- *Access network authentication* is necessary to provide the employee/user with network access. This authentication function must be supported with a wide variety of access networks and technologies that may include cellular service across multiple operators, Wi-Fi access in both internal and external (public hotspots) environments, as well as other network technologies (e.g., private LTE or 5G, CBRS and custom IoT network systems). In addition, VPN credentials may be required to access the private enterprise network from a public network. Each of these systems has different authentication mechanisms and may also have separate charging sytems that impact network selection decisions.
- *Application authentication* is necessary so employees/users have access to the applications needed in the day-to-day operation of the business. These applications can range from dedicated enterprise applications that run in a protected enterprise network environment to Office 365 tools that may run as SaaS in external cloud environments. Authentication mechanisms for each application may be different and also may have authorization and charging requirements that must be managed. In many cases, context awareness will be an important factor to support multi-factor authentication for these applications.

Generally, identity management solutions address parts of these authentication scenarios and may introduce new actors into the system with a new and different set of identity providers. Very large enterprises may have the IT scale to create besoke solutions to meet the above needs for at least their “main line” networks and applications. However, the ability to manage credentials, policies and authorization types across this complex environment is not easily available to most organizations, particularly small and medium enterprises.

3.2.1 Story Highlights

A small-enterprise employee is responsible for managing an aspect of the business where travel to other branch sites is required. Additional requirements include network access at each site, remote access while en route, application access within a branch or main office or remotely from a public network and access to a variety of cloud-based services for employee scheduling, product management, customer management and other business critical systems. To optimize efficiency and security, it is desirable that the employee can access these networks and applications from multiple devices using one master credential/identity management framework, which can set appropriate authorization and charging policies as needed.

3.2.2 Business Drivers

A disparate, complex set of authentication/authorization instances can have a significant business impact. For example, this environment can result in:

- Security-related events due to "shortcuts" made by employees that jeopardize security. This would include ineffective passwords that are easy to remember, and other poor password-related practices.
- Time and efficiency loss dealing with problems accessing the needed networks and applications to do the work required.
- Added support costs in debugging a complex environment where employees are unable to access the networks and tools needed.

3.2.3 Actors

Actors can include:

- The enterprise with enterprise networks and applications.
- Third-party networks and applications that have a business relationship with the enterprise, such as network service providers.
- Public networks and providers for all types of network access.
- Identity providers.

3.2.4 Deployment Model and High-Level Architectural Context

Solutions covering this case may require a framework to enable the management of policy and charging across a wide range of networks and applications.

These solutions should address the need for secure application authentication across the broad ecosystem of enterprise dedicated and outsourced applications. They also should support access network identify federation mechanisms to integrate CBRS and private LTE credentials with enterprise and/or mobile operator authentication.

3.3 MULTIPATH

Many devices support multiple network connections. For example, smartphones and many tablets connect not only to their associated mobile operator's network, but also support a Wi-Fi connection. In the future, these devices may also have additional network interfaces for private LTE, 5G and CBRS networks, as well as potential 6G nearfield technologies.

Although multipath standards and solutions have been available for many years, a consistent, easy-to-use platform that works across all applications and access networks remains elusive.

3.3.1 Story Highlights

An enterprise user has sufficient network connection speed at home or in the office. However, when traveling to meet customers, perform services or conduct other field-related business, high-speed access can be inconsistent or non-existent for any of the available network options. This lack of suffiently high bandwidth makes it difficult for the enterprise user to access the necessary enterprise applications and resources to efficiently perform their job. Access to an enterprise platform that enables secure, automatically enabled multipath access provides a more consistent user experience and enhances productivity.

3.3.2 Business Drivers

When critical enterprise applications can be accessed securely through a multipath platform that is "automatically" enabled, enterprise user productivity is enhanced. This can result in real, measureable productivity gains because enterprise users can work faster and more efficiently.

3.3.3 Actors

Actors can include:

- The enterprise with enterprise networks and applications.
- Third-party networks and applications with a business relationship with the enterprise, such as network service providers.
- Public networks and providers for all types of network access.

3.3.4 Deployment Model and High-Level Architectural Context

Many multipath protocols have been defined within the IETF. 3GPP has also defined architectural aspects to support multipath within the 5G core. A more detailed analysis of mulitpath is provided in the next section of this document.

4
ENABLING TECHNOLOGIES

4.1
POLICY PROPAGATION, DISCOVERY AND ENFORCEMENT

Policy is simply the set of rules used by an entity (app, enterprise, network operator) to control the behavior of a connection based on user identity, device and environmental factors (e.g., time of day, means of access, cost, access type).

Policy determines:

- Authentication: based on connection context.
- Authorization (segmentation): What is permitted to connect to what?
- QoS: setting the appropriate connection attributes (e.g., priority, bandwidth, routing).
- Security.
- Access and path selection based on:
  - Connection capabilities.
  - Application requirements.
  - Cost.

To deliver multi-network enterprise solutions, a client device must be able to connect to server infrastructure across a variety of access methods and transport networks. The client device might be a smartphone running a user application or any form of machine connection (e.g., IoT devices) that requires some form of (mostly) wireless connection.

Figure 4.1 illustrates the logical layers that make up this connection.

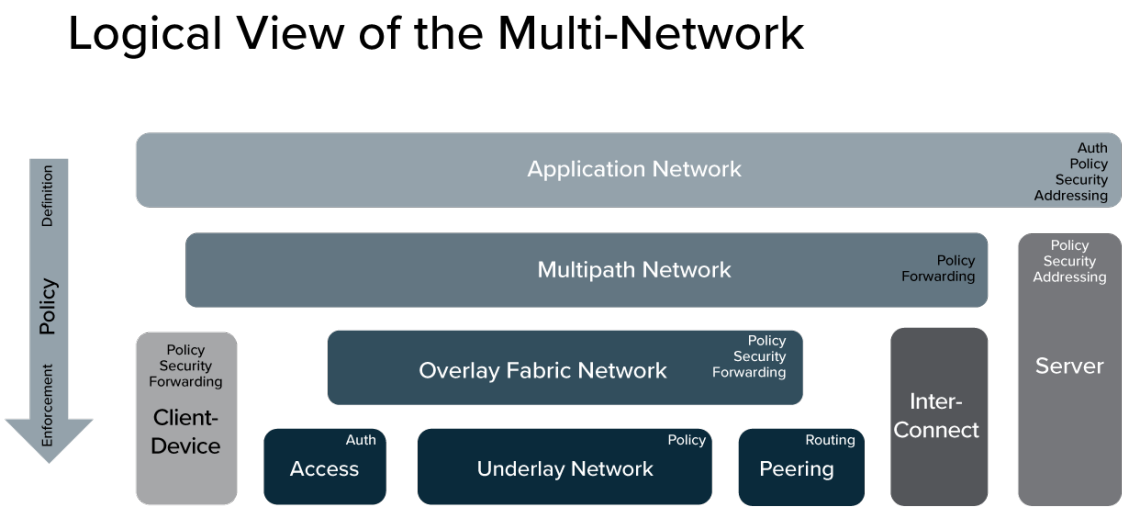

Figure 4.1 – Logical View of the Multi-Network

At the lowest level are one or more access methods, each with its own access control authentication method. This includes all methods of connectivity (e.g., Wi-Fi, cellular, LoRa, C-IOT and physical WAN for an enterprise network).

This model requires a physical underlay network to connect the access node to a gateway and ultimately a peering point. This includes backhaul and core for a cellular system, or Ethernet LAN and physical WAN for an enterprise network. The peering architecture is the physical provider transit interconnect where multiple peering methods can exist, ranging from settlement-free to paid private. Across the multitude of relationships and provider/user parties, the goal is always to transit between one underlay network to another.

In the figure, the service infrastructure is generalized into a single component because this report focuses on the connections and their characteristics. This includes data center (DCI, routing, switching), compute, storage and cloud (private or public).

The next layer up is the overlay network, which generalizes the service network from a service provider or the fabric network within an enterprise. This includes hub-and-spoke tunnels, mobile packet core forwarders, enterprise fabric tunnels, (SD-)WAN tunnels. This layer also generalizes the client addressing layer and the point where IP attachment is made.

Moving to the top of the figure, the application network is the way applications consider connectivity between the client and the server. Historically this occurred at layer 4 with socket connections. Although this is still the case for tightly coupled client to server connections, the application layer interaction between client and server now occurs abstracted from the connection method. Moreover, the usage of names for server element identification is the only visible identification mechanism, making the layer 3 and 4 addressing methods completely invisible.

Finally, the multipath network sits between the overlay network (service connection) and the application network (application procedures). Appreciating the application network abstraction and the need to connect across any access method and service connection, the multipath network enables a more seamless connection. It is here that multipath protocols (e.g., hICN, MP-TCP, MP-QUIC) are introduced to meet the needs of the client-to-server connection.

Figure 4.1 also exposes the methods by which traffic policies are considered and distributed. The highest and most abstract layer makes policy definitions — with little opportunity to enforce them. Moving down the layers, there are increasing capabilities and opportunities to enforce the application policy definition. Moreover, the lower layers, particularly as they are embodied in different provider/enterprise networks, may have specific policies of their own that they define and wish to enforce. Figure 4.2 illustrates the abstracted network with policy propagation shown.

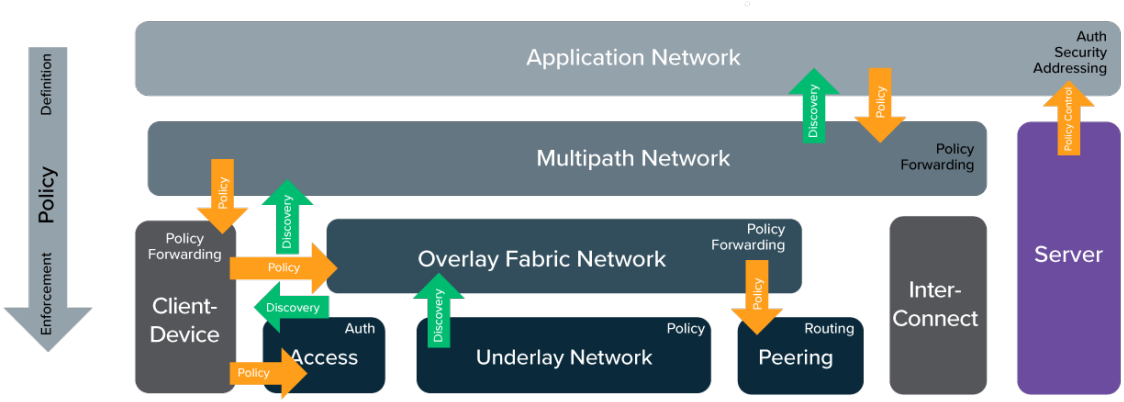

Figure 4.2 – Network Policy

The various components of the abstracted network and their roles in policy definition, propagation and enforcement are:

- **Client Device:**
  - Policy enforcement for traffic flows and multipath access methods.
  - Security functionality for root of trust and identiy verification.
  - Forwarding function for multipath policy execution.
- **Access**
  - Authentication and authorization point to use the accesss domain.
  - Multiple access methods possible, each with its own authentication framework.
- **Underlay Network**
  - The physical infrastructure for a connection network between access and peering.
  - Transport network of a service provider or an enterprise LAN/WAN.
  - Policy enforcement at the transport level linked to access method and traffic marking.
- **Peering and Interconnect**
  - Physical provider peering and interconnect, including provider backbone to connect between underlay network to peering infrastructure
- **Server**
  - Generalized data center or cloud infrastructure inclusive of DCI, routing, switching, compute and storage, public or private.
  - Policy enforcement in and out of the server environment.
  - Physical, gateway and analytical security.
  - Addressable at both the IP and application level.
- **Overlay Fabric**
  - The network overlay that the client devices operate across, such as hub/spoke tunnels, mobility forwarders, fabric tunnels, (SD-)WAN tunnels.
  - Policy enforcement at the forwarding level with traffic tagging.
  - Security for network utilization based on client device and user verification.
  - Forwarding layer as a continuous overlay between lower layer access and peering.
- **Multipath Network**
  - Above all access and forwarding overlay networks, client devices connect to servers across an any-path/any-access method.
  - This layer enables the continuity of application experience irrespective of the transitions across multiple access methods.
  - Policy enforcement for the multipath forwarding functions from client device to server through multipath forwarders.
- **Application Network**
  - Client application to application server as a communication network fully abstracted from the network(s) beneath.
  - Usage authentication and secure utilization through application identities and E2E encryption.
  - Addressing through namespaces, with no need for IP address continuity or even any IP-level visibility.

Figure 4.3 shows how some of the functions and capabilities required by multi-network enterprise solutions are mapped to the logical multi-network.

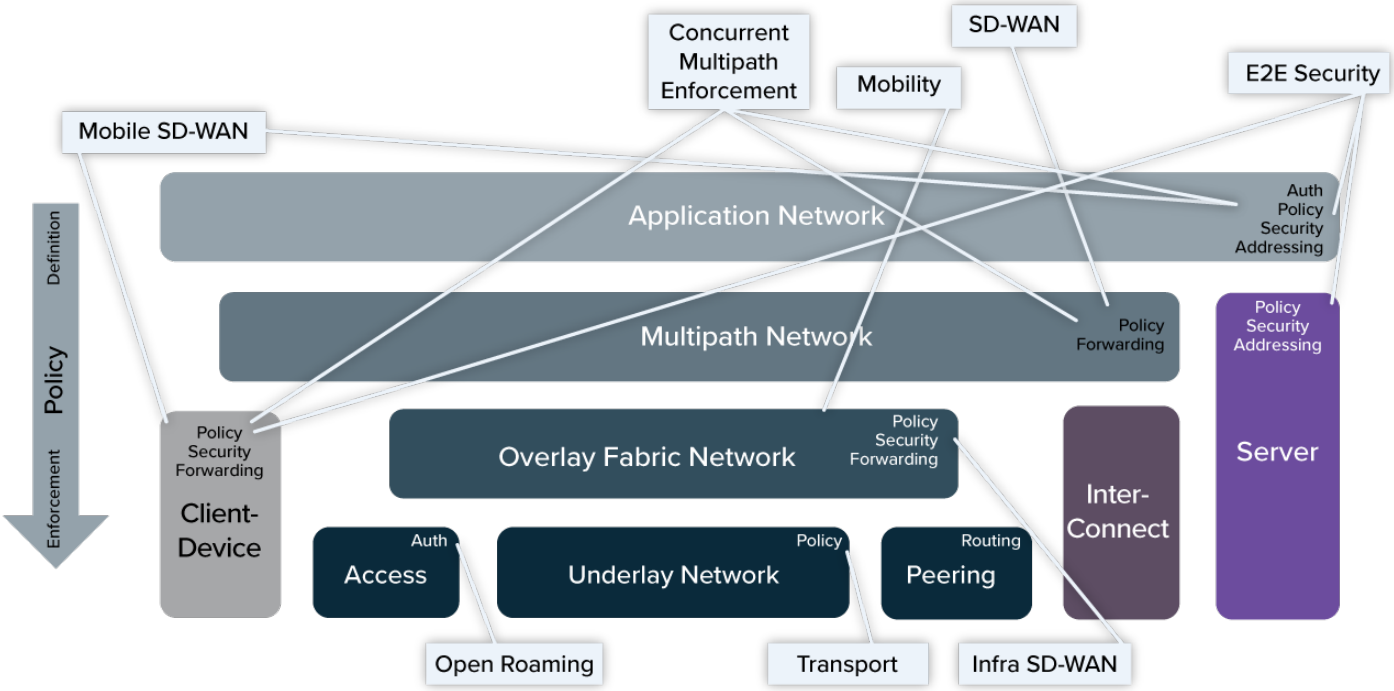

Figure 4.3 – Logical Network View – Functions

Management of a comprehensive E2E security strategy requires security policy propagation throughout the system. For example, zero trust architectures may require a PKI and a centralized authentication and authorization policy engine linked to system policy. E2E system policies would be used to manage security for data at rest, data in motion and data in use.

Data-in-use security is enabled through the use of Confidential Computing (e.g., using a Trusted Execution Environment). This capability has emerged through new, CPU-based technology that secures data at runtime. The technology is commonly available in various CPU architectures. Integrating the confidential computing capable silicon with zero trust based network/service access software will be highly desired.

Other functions noted above are described in more detail in the following sections.

4.2 IDENTITY FEDERATION

In the context of this report, identity federation is the process of delegating an individual's or entity's authentication and/or authorization decision to a trusted external party. Each party in federation plays the role of either an identity provider (IdP) or a service provider (SP) for the user/client. In identity federation, an IdP vouches for the identity of the users, while an SP provides services and access to resources to the users.

Identity federation can be used for network access or application access. For example, in an enterprise scenario, identity federation could be used to:

- Enable an enterprise user who is roaming outside the enterprise network to access a third-party Wi-Fi service using enterprise credentials (enterprise as an IdP).
- Enable an enterprise user to use a third-party enterprise application that may be hosted in the cloud or other third-party data center using Eenterprise credentials (enterprise as an IdP).

4.2.1 Access Identity Federation

A well-known, mature example of access identity federation is the 3GPP roaming model. 3GPP based devices can roam globally across hundreds of different mobile network providers in a relatively seamless manner.

An excellent example of non-3GPP access based identity federation can be found in WBA's OpenRoaming™ project. [5] It defines a cloud-based roaming federation that leverages native Passpoint® device capability, a PKI and a common legal framework to significantly lower the barriers to adoption of roaming services by all Wi-Fi networks.

OpenRoaming encompasses four key elements:

- *Cloud federation* creates a federation of access networks and identity providers to enable automatic roaming and user onboarding. The federation is based on WBA's Wireless Roaming Intermediary eXchange (WRIX) standards to scale and facilitate different business models under a harmonized framework. Minimum service level requirements for all OpenRoaming access networks ensure a baseline quality of service (QoS) is provided to all federation users.
- *Cyber Security* enables simple, secure and scalable exchange of roaming capabilities amongst different organizations that are part of WBA OpenRoaming. Allowing automatic and secure roaming between millions of networks, nationally and globally with secured interconnection and encrypted communications that are used to authenticate user devices using 3GPP-defined authentication exchanges.
- *Legal Framework:* To avoid the intrusive requirement for end users to endlessly agree to the separate terms and conditions on each and every network, OpenRoaming is built on a foundation of baseline legal agreements. The federation defines the set of baseline privacy policies and end-user terms and conditions that are required to be supported by all federation providers. Templated terms are defined that are used between OpenRoaming brokers and their provider customers to ensure all participants adhere to the OpenRoaming specifications.

5 https://wballiance.com/openroaming/

- *Network automation* defines the use of a roaming consortium codes framework (RCOI) to support policy provision on wireless devices and networks. RCOIs allow the policies of individual access and identity providers to differ. They also enable an automated matching of policies where authentication capabilities with a device are only triggered if there is a policy in common between access and the device's identity provider.

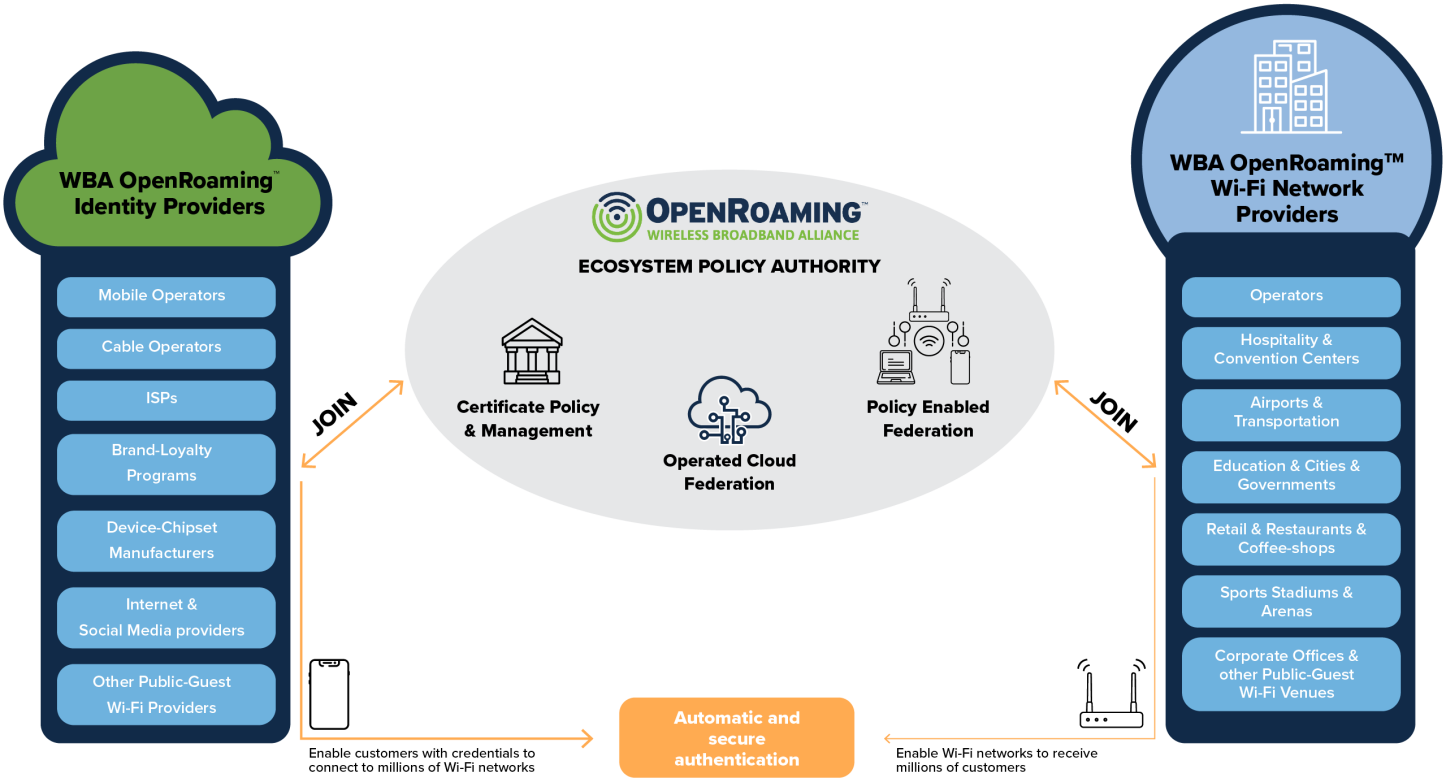

Figure 4.4 – OpenRoaming[6]

In the context of multi-network enterprise solutions, the WBA OpenRoaming system would enable the enterprise to register directly (or use a third party) as an identity provider and then enable enterprise employees and/or their customers to connect to the large set of OpenRoaming providers.

4.2.2 Application Identity Federation

When a user attempts to log into a specific application, the application can communicate with the IdP (facilitated by the user) to authenticate the user. The actual federation is often executed through standards such as OpenID Connect, Oauth or Security Assertion Markup Language (SAML).

OpenID specifications are promoted by the non-profit OpenID Foundation. [7] Its authentication protocols include OpenID 1.0, OpenID 2.0 and OpenID Connect. OpenID Connect 1.0 is a simple identity layer on top of the Oauth 2.0 framework.

Oauth 2 was developed within the IETF Oauth Working Group with the Oauth 2.0 Authorization Framework [RFC6749] and Oauth 2.0 Bearer Token Usage [RFC6750] specifications. These specifications provide a general framework for third-party applications to obtain and use specific HTTP resources. Specifically, they define mechanisms to obtain and use access tokens to access resources but do not define standard methods to provide identity information.

6 https://wballiance.com/openroaming/how-it-works/

7 OpenID Foundatation - https://openid.net/foundation/

OpenID Connect is an authentication protocol, while Oauth is for authorization. However, because OpenID Connect piggybacks on top of Oauth2, Oauth2 values can exist alongside OpenID Connect values in both the request parameters and the response payloads (when the APIs allow).

SAML is an open standard that allows identity providers to pass credentials to an application server to enable user authentication and authorization. SAML uses Extensible Markup Language (XML) with certificates instead of the access tokens more characteristic of OpenID Connect.

OpenID Flow (example)

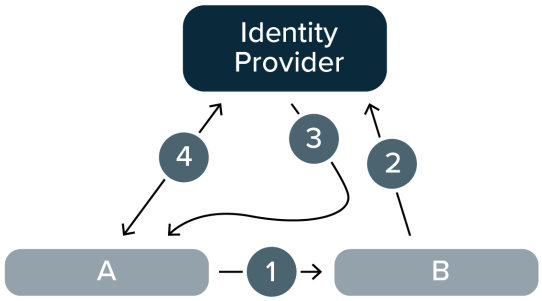

A: Application that wants to authenticate a user (B)
B: User

1 Application A needs to verify the identity of User B and first redirects B to the Identity Provider with the client ID of Application A.

2 User B signs into the Identity Provider and gets authorization to connect to the "client ID" (A).

3 Identity Provider redirects to the application (A) with either a signed token or authorization code.

4 Application A uses the authorization code to get an access token (or verifies token signature) for User B authentication.

Figure 4.6 – OpenID Example Flow

SAML Flow (example)

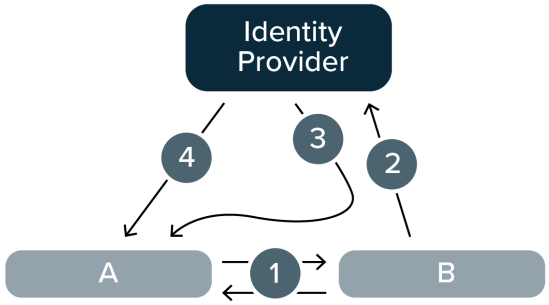

A: Application that wants to authenticate a user (B)
B: User

1 Application A needs to verify the identity of User B and creates a SAML request and redirects the User B to the Identity Provider with the request.

2 User B logs into the Identity Provider which builds the authentication response in the form of an XML document containing the user's identifier and signs it using an X.509 certificate.

3 Identity Provider posts this information to the Application A via a redirect.

4 Application A knows the Identity Provider public key and can validate the received certificate to enable User B authentication.

Figure 4.7– SAML Example Flow

These application-oriented identity federation standards work in very similar ways. They all enable a user to authenticate with an application server by leveraging the user's relationship with the IdP in a relatively seamless manner.

When the enterprise can provide the IdP function directly or via a third-party service, the enterprise can then manage user access to applications in one place, via control of the user's authorization profile as the IdP. In addition, control of the IdP function within the enterprise enables the enterprise to better control user access policies and the authorization of information exchange in a consistent manner.

4.3 DLT-ENABLED SELF-SOVEREIGN IDENTITY

Self-sovereign identity starts with the notion that we all are the makers of our own identity, online and off. Because they do not rely on any centralized authority, self-sovereign identity systems are decentralized, mirroring the way identity works in real life. Our interactions flexibly support the use of attributes and credentials from numerous third parties, all presented by the person they're about, typically by taking those credentials out of a wallet or purse and presenting them to someone else to verify.

Self-sovereign identity gives people control over their personal information, including the mechanisms to grant, revoke and apply constraints to their consent for others to use it. Self-sovereign identity makes use of decentralized identifiers (DID) and their credentials that are granted by an "issuer" to, and conveyed by, the user identity "subject" they are about. The issuer third-party credentials claims can be verified as authentic even when they're conveyed by the subject of the claim. These "verifiable claims" are the heart of self-sovereign identity.

Technological advancements now include ways to store, share and authenticate identity data and third-party credentials in a secure and privacy-preserving manner. Self-soverign identity entrusts an individual with ownership and control of their identity without interventions by administrative authorities or centralized or federated IdPs. This new ecosystem lets individuals interact directly with entities without the need for intermediaries, using public-key cryptography, decentralized identifiers and blockchain technology.

4.3.1 How DLT/Blockchain Platforms Can Solve the Identity Dilemma

Self-sovereign identity is a new approach to digital identity that gives individuals control of their digital identities. Self-sovereign identity solutions are being applied to a variety of domains where a new approach to identity is advantageous. This new technology may benefit enterprise authentication and authorization solutions and is an area of active study.

Self-sovereign identity systems use distributed ledger technology (DLT)/blockchain so that decentralized identifiers can be searched without involving a central directory. Blockchains don't solve the identity problem by themselves, but they do provide a missing link that allows things we've known about cryptography for decades to suddenly be used. That allows people to prove things about themselves using decentralized, verifiable credentials just as they do offline.

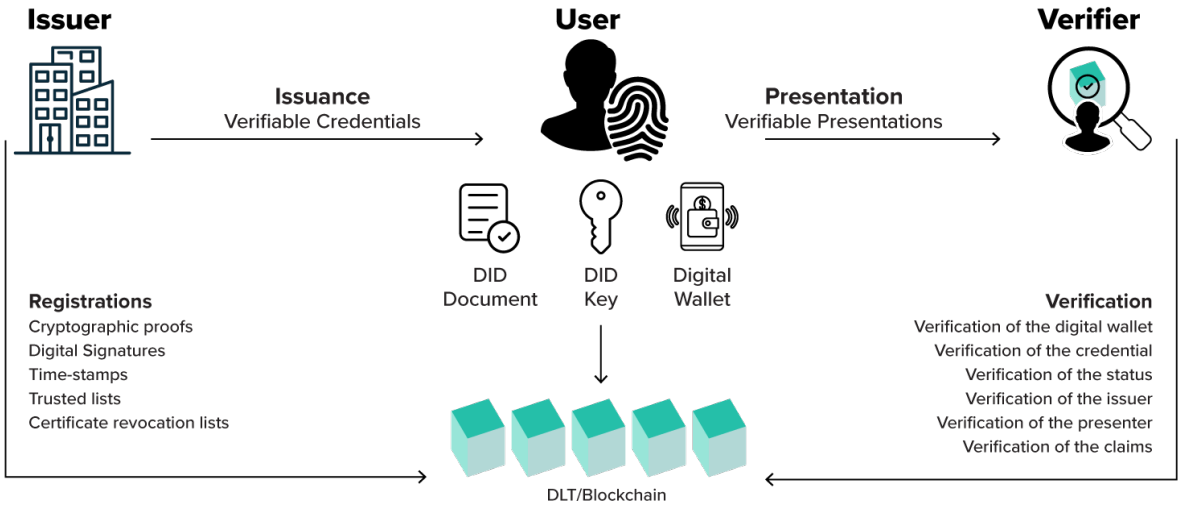

Figure 4.8 – DLT-Based Self-Sovereign Identity

MNES Using Self-Sovereign Identity

Enabling multi-network enterprise solutions using self-sovereign identity would enable an enterprise to issue verifiable credentials to an employee identity. This would provide their connected device(s) and applications with seamless connectivity to networks and enterprise services that are fully under the enterprise's control.

Suppose that the enterprise wants to change the permissions and policies or even revoke access for the employee, their device or application. It can simply update the verrifable credential for the employee identity.

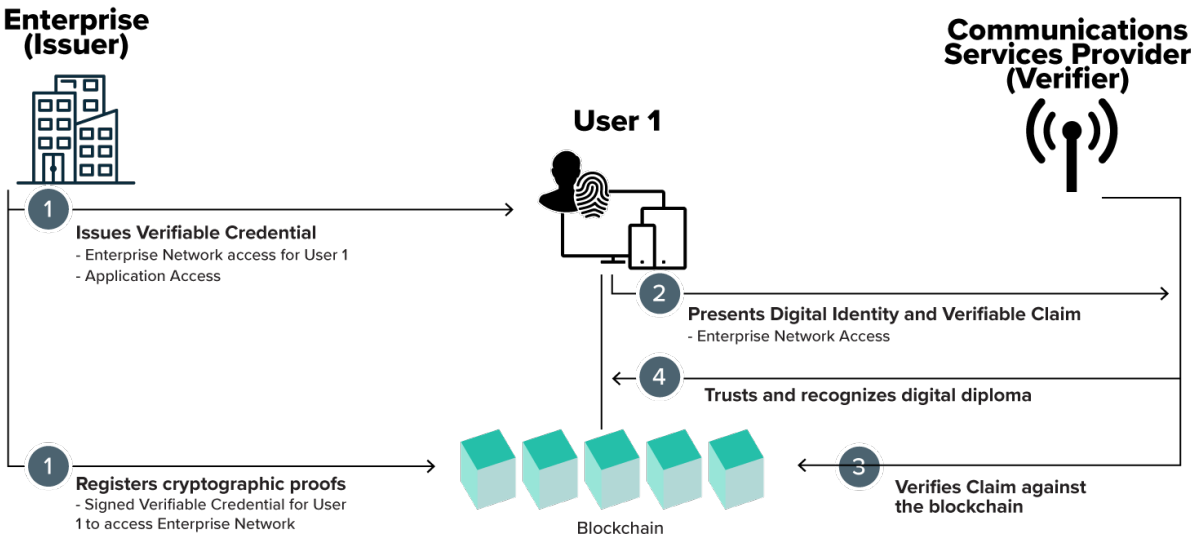

Figure 4.9 – Blockchain-Based Self-Sovereign Identity

4.4 MULTIPATH NETWORKING

Establishing and being able to use multiple paths enhances a session's robustness. It allows the best path to be used at all times, even if only one of the paths is active at any given time. If more than one path can be used to transmit data simultaneously, the throughput experienced by the application can be increased by permitting all available radio resources and access networks to be used for the session.

Enterprises and verticals will necessarily make use of multi-access, cross-sector, cross-domain (i.e., multipath) connectivity to enhance robustness and session integrity, as well as to increase throughput. In the multi-network scenario, the enterprise application may make use of both 3GPP-compliant and non-3GPP-compliant access networks, public and private networks and cloud providers. Enterprise multipath value propositions include:

- **Increased usage of all access networks:**
  - An enterprise application could use multipath to enable overlay and ad-hoc network utilization where available without user interaction.

- **Improved application quality of experience:**
  - Greater probability of being able to deliver critical applications, taking advantage of all available connectivity options.
  - Real-time services could be transported preferentially over access networks that have differentiated QoS enabled.
  - Interactive services could be transported using the lowest latency network available.
  - Bulk data services could be transported using the aggregated throughput of multiple different access networks.
- **Improved service resilience and enhanced robustness:**
  - Seamless switching from one access network (e.g., one with indoor coverage) to another (e.g., with outdoor coverage).
- **Lowering communications costs:**
  - Non-business-relevant applications would use a high-tariff access network only when lower cost alternatives are unavailable.
  - A service provider could manage resource utilization by switching traffic paths to control load in real time.

The particular multipath policies and mechanisms that should be applied can be defined by the application itself, the enterprise providing the application and by the access network operator.

Figure 4.10 illustrates three high-level models for multipath operation.

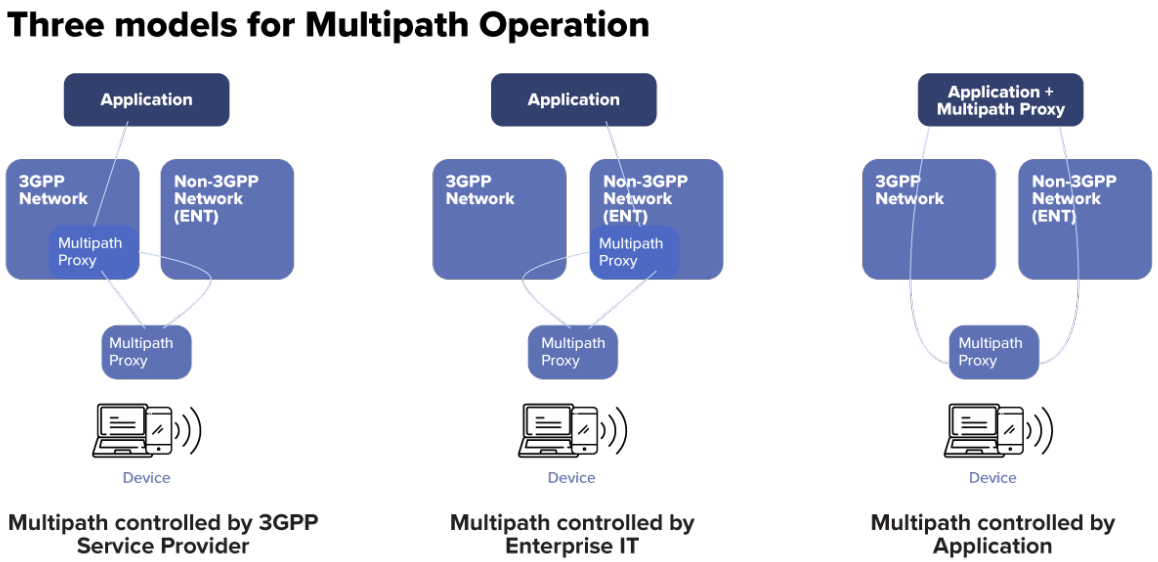

Figure 4.10 illustrates three high-level models for multipath operation.

All three of these models are important for MNES, although each is optimal for a different set of objectives.

4.4.1 Multipath Under Application Control

In this model, the multipath functionality is either intrinsically supported by the application or is implemented in a proxy above any transport network. The application is in the best position to understand its needs and which capabilities it can take advantage of. For example, an application knows how long sockets are likely to persist and can use that knowledge to determine whether to establish multipath operation.

In the ideal scenario, the application and all the intervening networks between the application server and the user device cooperate to realize the best multipath policies to support that application's needs. This approach is already supported in device ecosystems (e.g., using the iOS URLSession API). It is possible to enable the application to set the "mpathservice type," which will convey the application's multipath policy to the intervening network(s).

The characteristics of this model are:

- Application aware. The networks below the application are not assumed to have any special capabilities. Real-time performance measurements are made by the application, and traffic steering is done above either transport network.
- Any protocol can be supported. Applications that are multipath aware can be designed with the multipath protocol in mind. Examples include Apple's Siri (MP-TCP) and Google's YouTube (QUIC).

This model's primary objective is to maintain and optimize the performance of multipath-aware applications over multiple networks to achieve a uniform, satisfactory user experience for that application. Hence this model is ideal for steering multipath traffic optimally from the application's point of view.

4.4.2 Multipath Controlled by the Enterprise IT

In some scenarios, either the enterprise or a network operator wants to determine multipath policy according to its own needs (e.g., for load balancing, security, priority based on user or application). The majority of applications might not be enhanced to become multipath aware, so there is a need for multipath proxies on both the client and the server side to deliver multipath functionality for unmodified applications.

The characteristics of this model are:

- The multipath proxy is located in the enterprise network and is under the control of the enterprise network. The 3GPP network is seen as another transport leg available in the session.
- Support for some application awareness. Enterprise policy can be applied based on the application, upon the user, etc.
- Any multipath protocol (e.g., MP-TCP, MP-QUIC, hybrid ICN) can be supported. The choice can be different for different flows/session, and can be based on performance, application-specific policy, etc.

Since multipath behavior is determined by the enterprise policies, this model is ideal for steering multipath traffic optimally from the enterprise point of view.

4.4.3 Multipath Controlled by 3GPP Service Provider

In this model, the multipath proxy is located in the 3GPP network's 5G core. The connection to the non-3GPP access network is via the Non 3GPP Interworking Function (N3IWF), whose role is to create a tunnel from the user device to the 5G core through the non-3GPP access network. As a result, the traffic through this tunnel appears to have originated via a 3GPP network connection. In 3GPP Release 15, the multipath proxy was located in the User Plane Function (UPF), which established a multipath connection between the UPF and the client's multipath proxy using the Multipath TCP.

In Release 15, the characteristics of this model are:

- No application awareness.
- The only multipath protocol supported is Multipath TCP.

An enhanced multipath functionality — Access Traffic Steering, Switch and Split (ATSSS) — is being standardized beginning with 3GPP Release 16 and extending into Release 17. Figure 4.11 illustrates the ATSSS architecture.

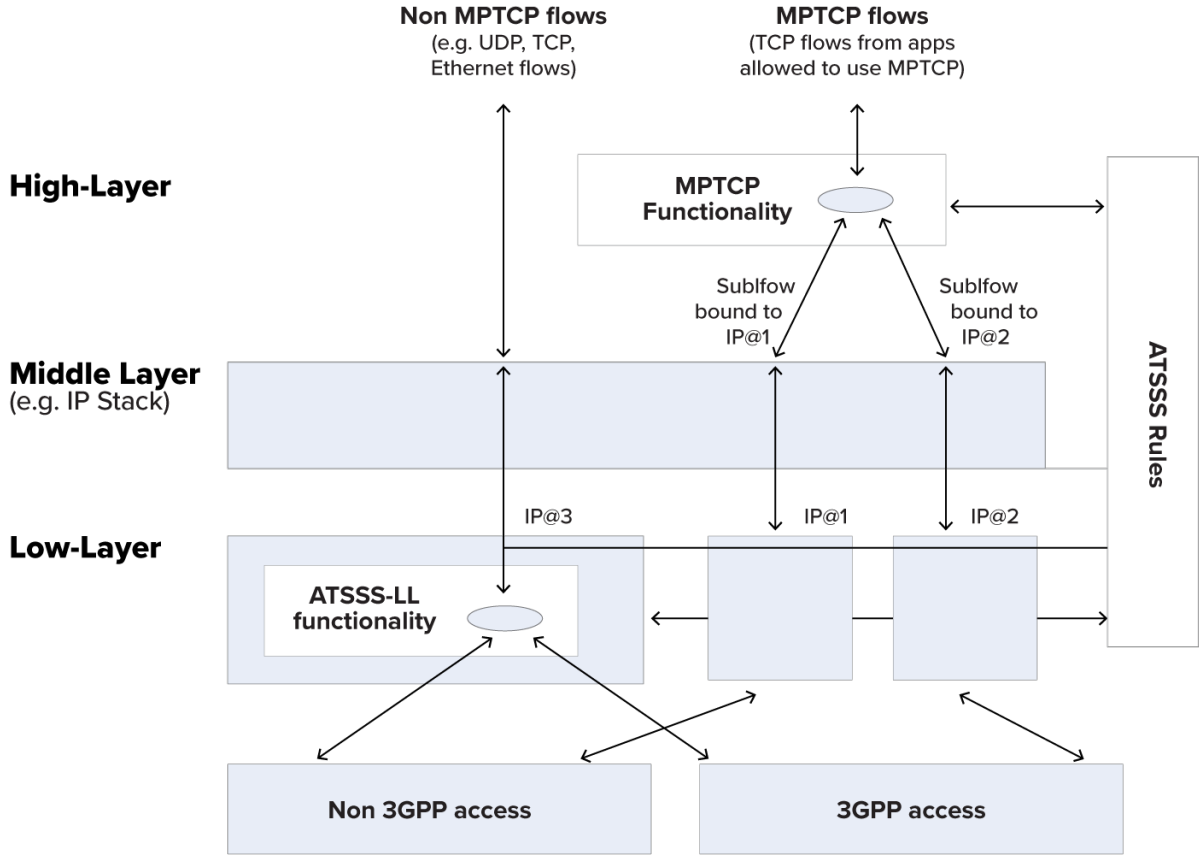

Figure 4.11 – ATSSS Architecture

Access Traffic Steering, Switch and Split (ATSSS)

ATSSS provides support for a multi-access PDU session. Support has been included for two steering functionalities:

- MP-TCP functionality for TCP traffic, with the MP-TCP proxy located in the UPF, using the MP-TCP protocol over the the 3GPP and/or the non-3GPP user plane.
- ATSSS-LL functionality for all traffic types, including TCP, UDP and Ethernet.

ATSSS includes provisions for initiating multi-access sessions (by the user device) and for storing and implementing rules governing them. The rules are constructed for the session by the Policy Control Function (PCF) in the 5G core, combining the user request with network operator policies, and then transmitted to the Session Management Function (SMF), which then stores the rules and transmits them to the user device and the user plane in the network.

There are several steering modes supported in ATSSS, each intended to optimize some aspect of the session:

- **Active-Standby:** This is used to steer traffic to one access link (the Active Access) when this access is available, and to switch the traffic to the other access (the Standby Access) when the Active Access becomes unavailable.
- **Smallest Delay:** This is used to steer traffic to the access that is determined to have the smallest round trip time (RTT) as measured by a new function, the Performance Measurement Function (PMF).
- **Load-Balancing:** This determines the percentage of the Service Data Flow (SDF) (i.e., the multipath session) that should be sent over the 3GPP access and over the non-3GPP access.
- **Priority Based:** This is used to steer all of an SDF's traffic to the high-priority access until this access is determined to be congested or becomes unavailable, in which case all traffic is switched to the low-priority access. This is a superset of the Active-Standby Mode.

4.5 SD-WAN PLATFORMS

ATSSS is somewhat more application-aware than the Release 15 multipath functionality in the sense that the user device signals the desire to initiate a multipath session. But there is still no intrinsic knowledge of the application's needs built into this model, nor any good mechanism for building awareness of enterprise needs into policies used to direct traffic. However, it is very well adapted to shaping multipath sessions optimally from the service provider's point of view.

SD-WAN refers to software-defined networking (SDN) in a wide-area network (WAN) context. SD-WAN has become a popular construct to support enterprise WAN connectivity needs because it leverages cloud-software-controlled management mechanisms to simplifiy a WAN's management and operation, decoupling the networking hardware from its control mechanism.

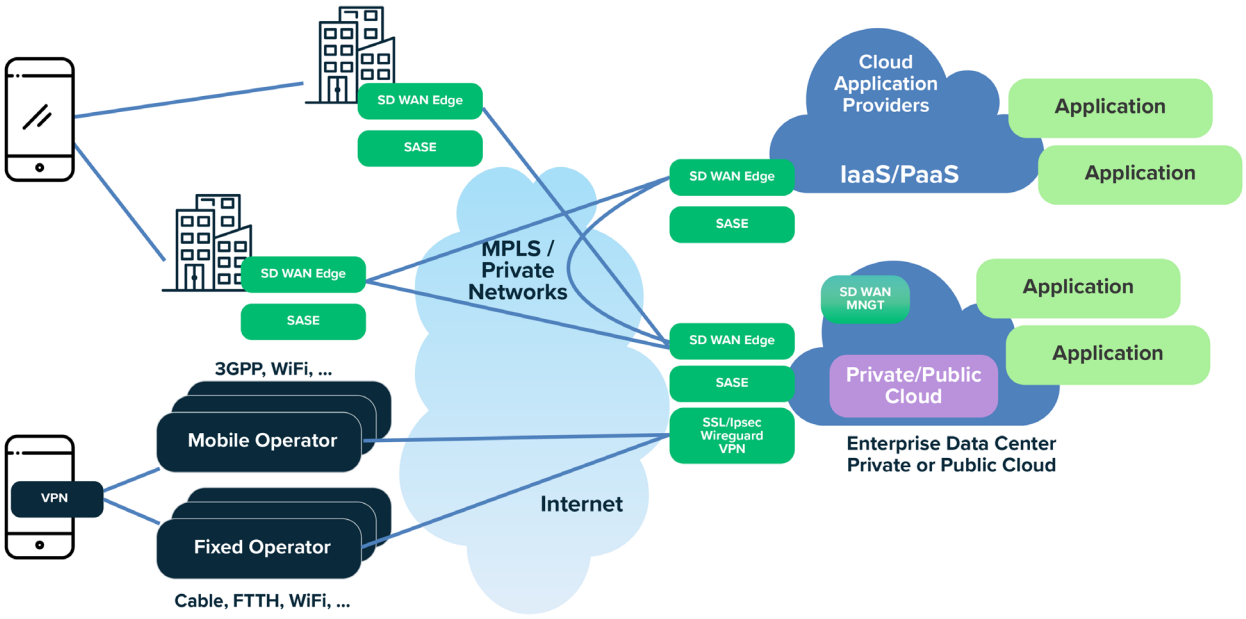

Figure 4.12 – SD-WAN Architecture

An SD-WAN minimally consists of an SD-WAN edge element and a centralized management function to provide SD-WAN control and orchestration. The SD-WAN edge is a software-controlled physical and/or virtual network function placed at the enterprise branch/regional/central office site, data center and in public or private cloud platforms. SD-WAN management includes orchestration and control functions used to configure the SD-WAN and set centralized policies used to make forwarding decisions for application flows.

An SD-WAN runs as an overlay on top of multiple types of network infrastructure from private/MPLS links to tunnels over the internet to network slices. SD-WAN edge elements can have multiple connections using different transport types, such as MPLS, internet, LTE and 5G.

A VPN may be used to connect to the SD-WAN, giving individual devices access to critical enterprise applications.

SASE recently emerged as a new network security service that enables flexible access from any network. SASE is a cloud-delivered service that identifies users and devices to apply policy-based security, delivering secure access to the appropriate application or data. The SD-WAN edge function is naturally located at a network edge and thus provides an ideal platform for SASE security capabilities. These capabilities may include a capabilities set of network security functions such as:

- Secure Web Gateway (SWG).
- Cloud Access Security Broker (CASB).
- Zero Trust Network Access (ZTNA).
- Remote browser isolation capabilities.
- Advanced threat analysis capabilities.

By placing a common SASE solution at the SD-WAN edge, a uniform level of security can be applied for application access from all potential access networks. In addition, security capabilities can be made aware of identity and context to make more informed, policy-based decisions regarding access and resource authorization.

Many SD-WAN platforms support a wide range of features including:

- The ability to pass operational rules and policies from the central SD-WAN controller to SD-WAN edge element.
- The ability to provide traffic shaping and QoS management functions based on these policies at the SD-WAN edge.
- Support for a variety of security features such as stateful application-aware firewalls, intrusion protection detection (IPS/IDS) systems, URL filtering for outgoing enterprise traffic and DNS security functions.

Some of the above functions may also be implemented by SASE together with SD-WAN edge elements for cloud-based services.

5
POTENTIAL SOLUTIONS

The enabling technologies discussed in the previous section of this document can be illustrated with the following example deployments. These examples are not intended to reflect real, deployable solutions. Instead, they place the enabling technology in a multi-network enterprise context to better illustrate how the enabling technologies may be used to address the challenges.

5.1 BUILDING ON SD-WAN PLATFORMS

The SD-WAN deployment can provide an effective platform for additional services such as zero trust network access, identity federation and multipath networking. It also can provide an infrastructure to enable policy propagation to better manage user device performance.

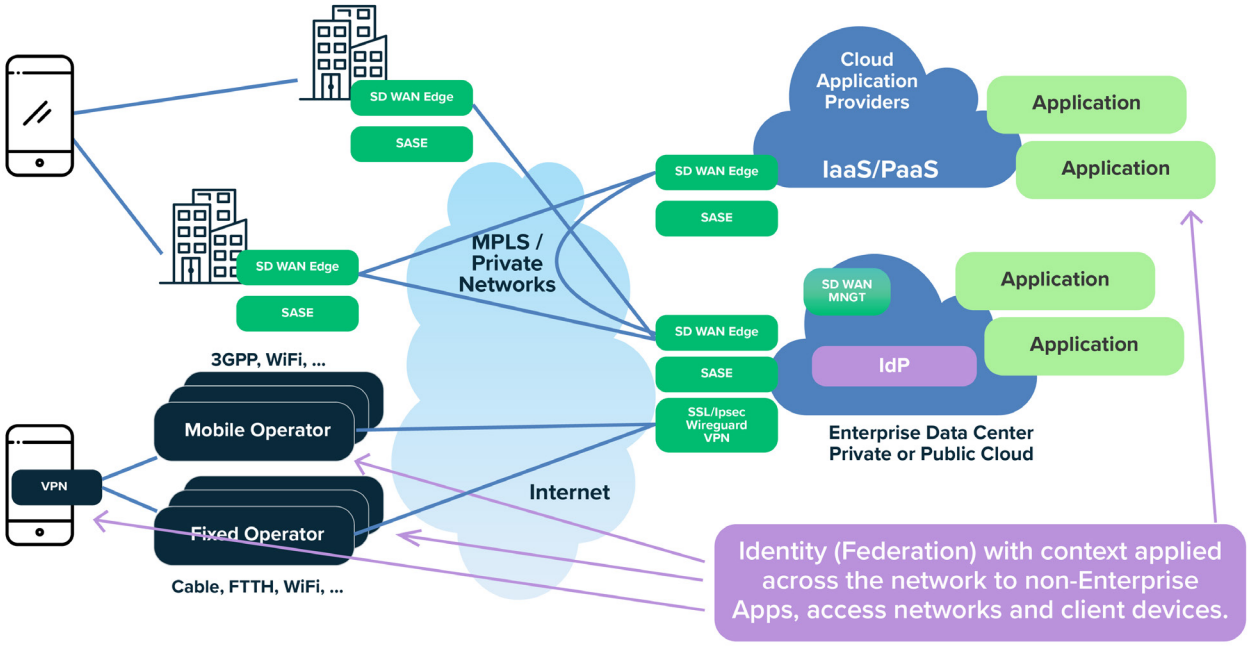

Figure 5.1 – SD-WAN with Added Identity Federation and Security Capabilities

An SD-WAN must be linked with an identity management function to authenticate and authorize enterprise users onto the SD-WAN itself, as well as to provide additional authorization for access to specific enterprise applications. For example, only selected employees who are involved with financial management of the enterprise should be authorized to access its financial applications. By expanding this identity management function into a more complete IdP function, both access- and application-federated identity services can be enabled. That is, the IdP function within the enterprise or in the cloud can then be used to authenticate and authorize enterprise usage of various access networks, as well as enterprise applications. One key advantage of this approach is that enterprise user authentication and authorization permissions can be managed from a “single pane of glass,” enabling fast and efficient employee onboarding and offboarding.

The IdP function is a seperable function in the enterprise, so it can can interwork with (or be outsourced to) external IdP companies such as a network service provider that could provide this function “as a service” to the enterprise. This lets the enterprise take advantage of application andaccess network federation/roaming arrangements created by the external IdP. For example, the external IdP may have already negotiated many agreements for both access providers (e.g., within the WBA Open Roaming Alliance) and application providers.

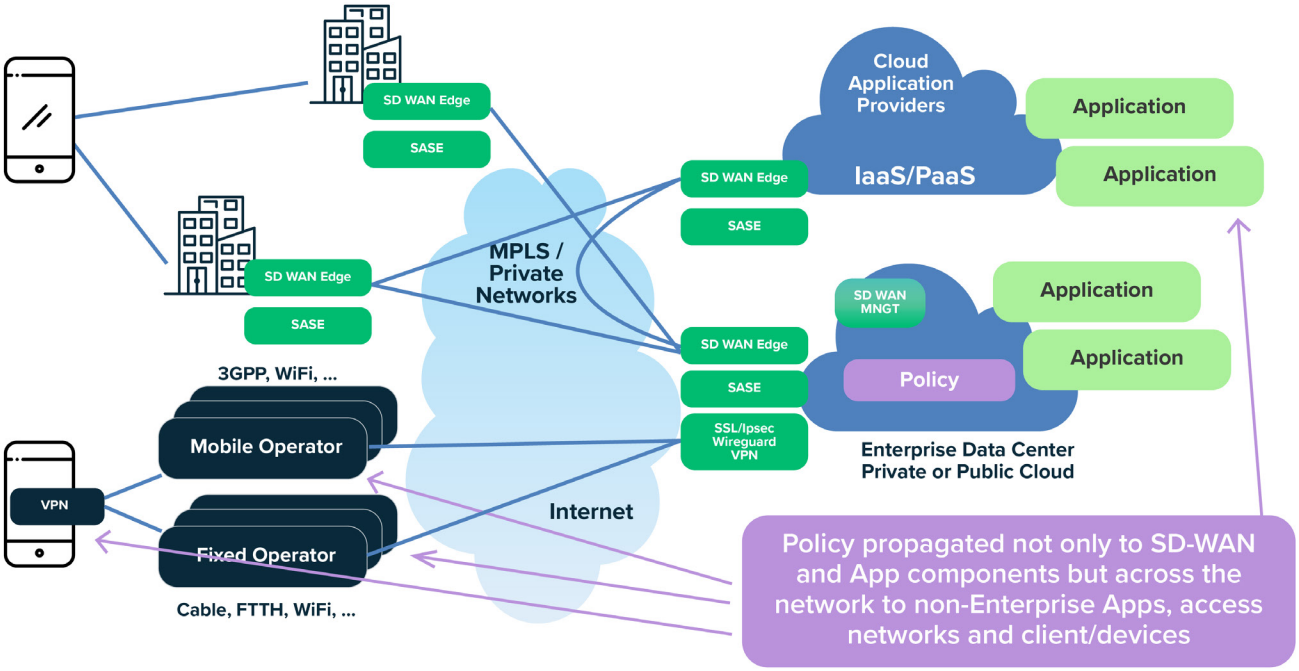

Figure 5.2 – SD-WAN with Enhanced Policy Capabilities

An SD-WAN solution also communicates network policies to the various SD-WAN edge elements in the SD-WAN private network. This existing network policy propagation function can then be leveraged for application- and/or QoS-based policies, as well as multipath policies when multipath services are deployed in the SD-WAN context.

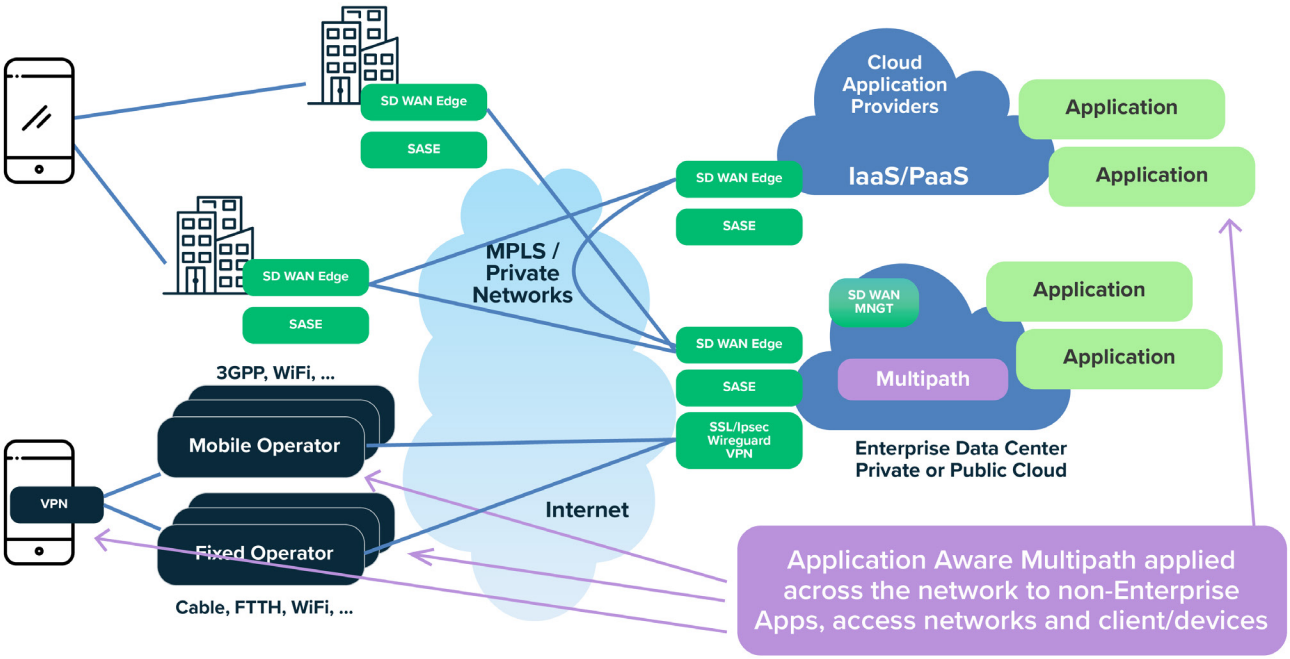

Figure 5.3 – SD-WAN with Added Multipath Capabilities

In situations where multipath capabilities are desired on a per-application basis, the multipath service node could be instantiated as a cloud function in line with the cloud application. This function could then be a separate function or integrated with the SD-WAN edge element that may already be present in the solution.

5.2 CLOUD-BASED PLATFORMS

Most of the enabling technologies described in Section 4.0 do not necessarily require a network-centric platform as described with SD-WAN in the previous section. In fact, key capabilities can be deployed across enterprise locations and applications by leveraging cloud-based services.

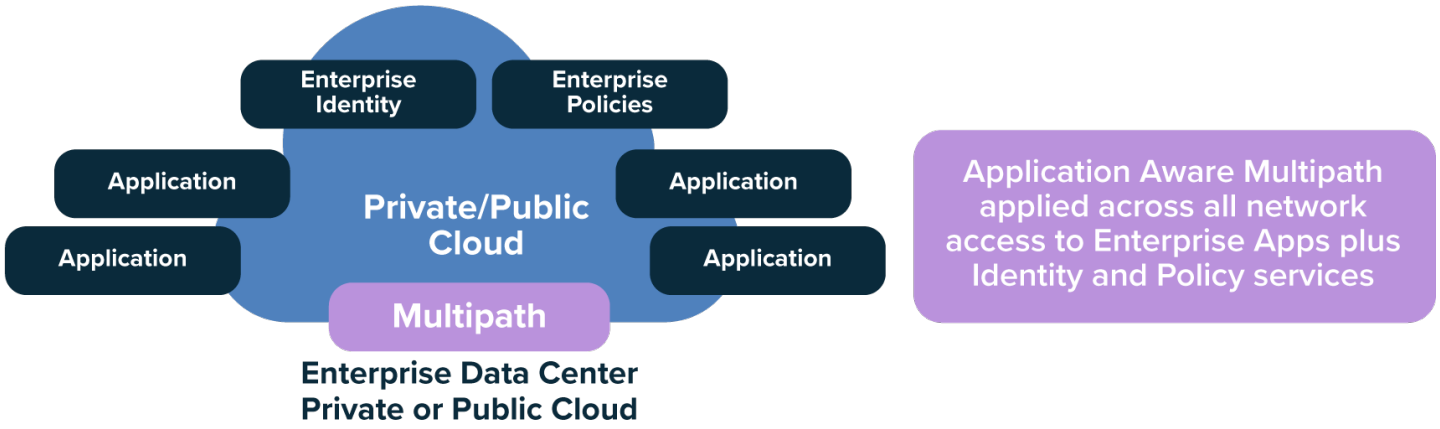

Figure 5.4 – Cloud-Based Platforms

Figure 5.4 shows an example cloud-based deployment that supports:

- A common federated identity service that could be used across enterprise applications for both application authentication and authorization, as well as network authentication and authorization (e.g., with WBA OpenRoaming).
- A common enterprise policy management service that could be leveraged by both the identity management function and the multipath service to propagate policies.
- A multipath capability that can be application oriented to provide better connection services for applications that can benefit from multipath networking.

A cloud-based platform can also be extended across multiple cloud application (IaaS/PaaS/SaaS) domains in providing multi-network services.

6
CONCLUSIONS

Enterprises utilize a wide variety of networks and network technologies to deliver services and support internal operations. However, controlling and managing this diverse communication environment is challenging because each network may have different authentication credentials with separate identity providers, as well as different security capabilities, policies, levels of performance and coverage domains.

This report introduced an industry approach to integrate and simplify the management of this complex network environment and enhance the overall reliability, resiliency, performance and security of the end-to-end system. Specifically, this report presented a set of enabling technologies that can be leveraged to address the problem space, including:

- Mechanisms to distribute a common, enterprise-controlled policy set to all network domains in a consistent manner.
- Common authentication and identity management structures to simplify authentication across all of the various network modes.
- Network capabilities such as multipath to better leverage multiple networks in concert to increase overall network reliability, resiliency and performance.
- Security capabilities to unify and enhance overall security across multiple networks and types. Specifically, security issues relating to PNT, the use of SASE to create a consistent security perimeter and the need for security for data-in-motion, data-at-rest and data-in-use.

These enabling technologies were further illustrated through example solutions leveraging both SD-WAN and cloud-based platforms. It is hoped that these technologies will enable enterprises to more efficiently manage their various networks, offering benefits to both the enterprises and their network providers. In this way, the enterprise achieves lower overall management cost while increasing security and performance.

BIBLIOGRAPHY

1. Say Hello to SASE (Secure Access Service Edge) (Gartner), https://blogs.gartner.com/andrew-lerner/2019/12/23/say-hello-sase-secure-access-service-edge/
2. Confidential Computing - Improve isolation of sensitive data payloads with hardware-based memory protections (Intel), https://www.intel.com/content/www/us/en/security/confidential-computing.html

ANNEX A

DEFINITIONS, ACRONYMS AND ABBREVIATIONS

Definitions

For a list of common communications terms and definitions, visit the ATIS Telecom Glossary, which is located at https://glossary.atis.org.

Acronyms and Abbreviations

ATIS	Alliance for Telecommunications Industry Solution
ATIS SYNC	ATIS Synchronization Committee develops standards pertaining to network synchronization
ATSSS	Access Traffic Steering, Switch and Split – standardized beginning with 3GPP Release 16
CASB	Cloud Access Security Broker
CBRS	Citizens Broadband Radio Service – a 150 MHz wide broadcast band of the 3.5 GHz band (3550 MHz to 3700 MHz)
DCI	Data Center Interconnect
DID	Decentralized Identifiers
DLT	Distributed Ledger Technology
DOCSIS	Data Over Cable Service Interface Specifications
FCC CFR	Federal Communications Commission Code of Federal Regulations
GPS	Global Positioning System
hICN	Hybrid ICN (Information Centric Networking)
IaaS	Infrastructure as a Service – Cloud model
IdP	Identity Provider
IPS/IDS	Intrusion Protection Detection
ITU-T	International Telecommunication Union – Telecommunications Standardization Sector
LORA	Refers to Long RAnge Wide Area (LPWA) networking protocol s designed to wirelessly interconnect Internet of Things (IoT)
MPLS	Multiprotocol Label Switching
MP-QUIC	Multi-Path – QUIC (a general-purpose transport layer network protocol that intends to improve performance of connection oriented applications)
MP-TCP	Multi-Path – Transport Control Protocol
N3IWF	Non 3GPP Interworking Function
PaaS	Platform as a Service – Cloud model
PCF	Policy Control Function
PKI	Public Key Infrastructure
PMF	Performance Measurement Function
PNT	Position Navigation and Timing
QoS	Quality of Service
RCOI	Roaming Consortium Codes Framework
RTT	Round Trip Time
SaaS	Software as a Service – Cloud model
SAML	Security Assertion Markup Language
SASE	Secure Access Service Edge
SDF	Service Data Flow
SD-WAN	Software Defined – Wide Area Network
SMF	Session Management Function
SP	Service Provider
SWG	Secure Web Gateway
UPF	User Plane Function
VPN	Virtual Private Network
WBA	Wireless Broadband Alliance
WRIX	Wireless Roaming Intermediary eXchange
ZTNA	Zero Trust Network Access