White paper: Hierarchical software defined networks for wide-area air traffic management networks

With new data hungry application flooding the air traffic management (ATM) world, driven by system-wide information management (SWIM), advanced aeronautical information management (AIM) data exchange, and the latest remote virtual tower (RVT) video applications, the ATM industry is looking at modern, safe, and secure network technologies such as software defined networks (SDN). Just a few years ago, only forward-leaning organisations embraced this compelling new concept to programmatically control their network resources.

Today, SDN is acknowledged as a major building block of next-generation ATM networks. The SDN concept separates the routing decisions, the 'control plane', from the actual data forwarding hardware, the 'data plane'. In SDN the routing decisions are performed in a separate SDN controller. Depending on the network architecture, SDN controllers can be deployed in a distributed or centralised fashion.

In the distributed approach, SDN controllers are deployed at various key locations within the network, each responsible for a sub-set of network devices with no communication to other SDN controllers. Although this concept allows for lean deployments at each site it may lead to unstable and/or asymmetric routing because of sub-optimal routing decisions by each individual controller due to the lack of information of the complete network.

To overcome this, in the centralised approach, clusters of synchronised SDN controllers are deployed at central locations within the network, each being able to take over the whole set of network devices. This allows for end-to-end optimised routing decisions as each cluster has all the network information available and the cluster members operate upon a synchronised information base.

A fault case analysis shows that the centralised approach may fail if the network becomes partitioned. Partitioning occurs if the SDN cluster loses connectivity to a portion of the network. These network elements are then considered orphaned. To ensure continuity of service, a hierarchical SDN deployment approach provides the best solution, where the central, synchronised SDN controllers are augmented by dedicated, distributed SDN controllers being able to take over orphaned devices.

This paper describes and analyses a hierarchical SDN architecture where a centralised SDN controller cluster manages the network during regular operation, and making use of a distributed SDN approach to ensure business continuity during degraded operation. Finally, network performance advantages of the hierarchical approach compared to conventional central SDN cluster deployments are highlighted.

Introduction

The transition of legacy systems into IP (i.e., TDM-over-IP towards VoIP), as well as the emergence of new data applications such as SWIM or RVT, are forcing air navigation service providers (ANSPs) to adopt converged, end-to-end all-IP infrastructure[1].

Compared to the enterprise telecommunication industry, an ATM-grade network defines a solution that fulfills certain safety-critical requirements applicable to the ATM context. Those requirements cover, for ground systems, aspects such as:

- Stringent quality of service (QoS) requirements: depending on the application there are strict requirements regarding delay, jitter or packet loss (i.e., ED-138/1).
- Reliability: depending on the application, the required end-to-end availability may be, at least, of five nines.
- Well-tested and documented: each component has to be developed according specific standards and guidelines (e.g., DO-278A/ED-109A or ED-153[2 3 4]) to allow a safety certification.

To support all these capabilities, operational automation is a major requirement to avoid slow and error-prone manual intervention. Thanks to the introduction of SDN technologies, the network is a completely flexible virtual entity defined via a software overlay. ATM-grade overlays sit between network infrastructure and applications, eliminating the strong dependencies that previously existed between individual applications and transport networks, whereas respecting application specifics. In this regard, the goal of this paper is to present different approaches for deployment of SDN-based solutions for ATM-grade networks.

The rest of the paper is organised as follows. First, our concept of an ATM-grade network is presented. Then, different SDN-based network deployments are described and analysed, with pros and cons. Finally, we discuss how our proposal may support the concept for aeronautical telecommunication network using internet protocol suite (ATN/IPS) standards and protocols[5].

Building an SDN-based ATM-grade network

To address the problems emerging in the transition to IP-based or hybrid converged networks, a new network architecture and design must be developed based on the following expectations:

- The network must be application-aware, with traffic management based on service quality and differentiated application profiles.
- Safety-critical methodology must be applied to network design, including resilience to multiple simultaneous failures and unpredictable anomalies.

Working from this premise, there are several distinguishing characteristics of an ATM-grade network that arise. Unlike a typical enterprise network, an ATM-grade network is:

- Deterministic in its performance.
- Resilient to unusual errors, including survivability to catastrophic event using satellite as backup, and security incidents, via proper network isolation and encryption.
- Built with the safety requirements of ATM applications in mind.
- Compliant with international rules and regulations for the ATM domain (i.e., ICAO, EUROCONTROL or RTCA).
- Operated according to well-defined and executed procedures, especially concerning changes and their effects.

To deliver ATM-grade networks, SDN technologies come into the market to provide the following capabilities:

- Scalable networks that integrate diverse technologies, including multi-vendor solutions.
- Situational awareness for end-to-end, real-time monitoring of all applications and networks.
- Intelligent routing and control to provide the performance that applications need.

The roots of SDN rely on the existence of an external "brain", the so-called SDN controller, outside the network equipment, which instructs them according to well-defined policies that cannot be implemented by the devices themselves because of limited intelligence or lack of global network view. However, this comes at the cost that a failure in the SDN controller, or in the communication channels towards the devices (referred to as control plane), cannot affect forwarding of the traffic (data plane). Therefore, this safety statement has two implications:

- Forwarding policies in the data plane shall (usually) remain in the last state before the control plane failure, until this one is recovered. While the last state may be not the current optimal at least all traffic is not dropped, and different degraded service profiles could be defined.

- Control plane traffic, that also goes through the data plane, shall be properly prioritised as non-best-effort traffic, maybe including dedicated bandwidth reservation.

As in enterprise telecommunication businesses, there are several approaches to provision an SDN solution towards ANSPs, which will be described in the following subsections. For the sake of simplicity, we are going to present SDN controllers in a very abstract way without any detail about real implementation. This detailed description is left to the next section.

Approach 1: SDN overlay

The basic offering for ATM-grade networks is based on the creation of an additional overlay on top of generic enterprise network services (hereafter, underlay networks). The overlay network is the integration layer between multiple underlay networks, presenting enhanced transport capabilities with improved safety (and security) parameters to the ATM applications.

One major advantage of this approach is that the ANSP can get an ATM-grade network availability from regular network offerings, as it can be demonstrated by following the well-known expression to compute availability for N parallel systems:

$$A = 1 - (1 - A1) \cdot \ldots \cdot (1 - AN) \tag{1}$$

In the previous example, assuming A1 = A2 = 0.99 (two-nines), then A = 0.9999 (four-nines). Interestingly, by adding a third two-nines operator, the overall availability grows to six-nines. However, this fact may be challenged in a real deployment, since it is not trivial, if even possible, to demonstrate that two operators are able to deliver an end-to-end physically disjoint network from one another for each site.

In this scenario, the SDN controller is usually referred to as SD-WAN (software-defined wide-area networking) controller. The reason is that for the controller the “network topology” is a set of sites with different applications feeding them with traffic (often receiving traffic as well) on the access side, and a set of “virtual” links (or transport pipes) through which they can reach the other sites in a single hop, because the underlay networks are transparent (see Figure 1a).

Figure 1: SDN approaches

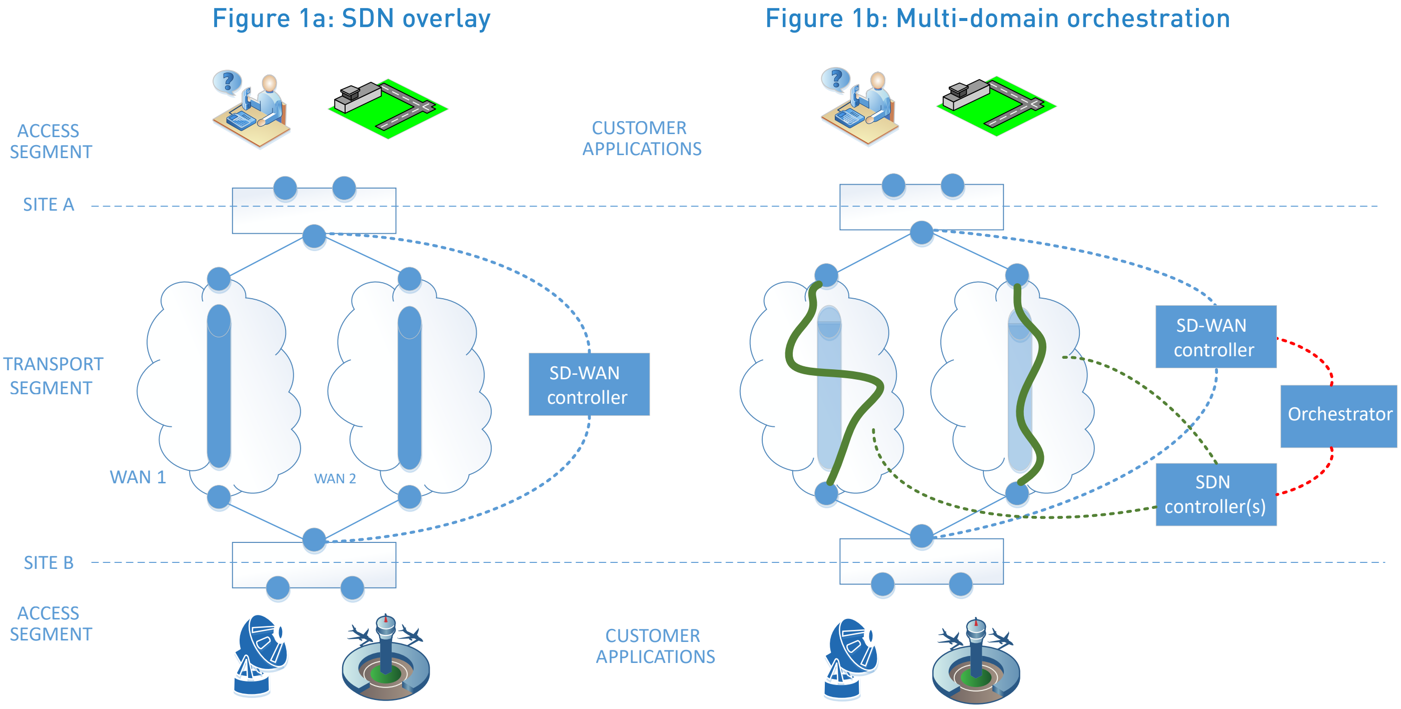

Figure 1a: SDN overlay

Figure 1b: Multi-domain orchestration

Clearly, the role of the SD-WAN controller is to select for each application the best WAN according to the network conditions.

To do this, SD-WAN controllers are complemented with capabilities such as active probing and passive monitoring.

Figure 2 presents the logical architecture of a customer premise equipment (CPE) from a generic site. On the access, different service delivery points (SDPs) allow the customer's applications to be attached to the WANs, for the transport. The CPE itself is composed of two items: (i) the policy-enforcement point (PEP); and (ii) probes.

On the one hand, the PEP is the network element through which the SD-WAN controller applies steering decisions. To this regard, it is out of the scope of the paper whether the WANs provide layer 2 or layer 3 transport services, and if steering in the CPE is performed in layer 2 or layer 3.

On the other hand, probes are the entities providing WAN monitoring capabilities. They are responsible for measuring the quality of each candidate WAN and deliver such results to the SD-WAN controller. In general, carrier-grade switches and routers (or PEPs in general) provide built-in probing capabilities, especially if functionalities like bidirectional forwarding detection[6] (BFD) are supported for link failure detection. However, it could be possible to deploy them as an external appliance.

It is worth mentioning that the SD-WAN solution for SDN overlays supports mixtures of heterogenous WAN connections, which means that it is expected to combine MPLS-based wired transport with wireless technologies like VSAT. In this regard, the SD-WAN solution is expected to provide extra functionalities to guarantee deterministic behaviors, for instance, to prioritise which WAN should be primarily used if all of them fulfill the QoS requirements and there is enough bandwidth, or apply admission control policies otherwise. In the following, we will assume that there is always enough bandwidth and steering is performed only according to probing results.

Figure 2: SD-WAN implementation

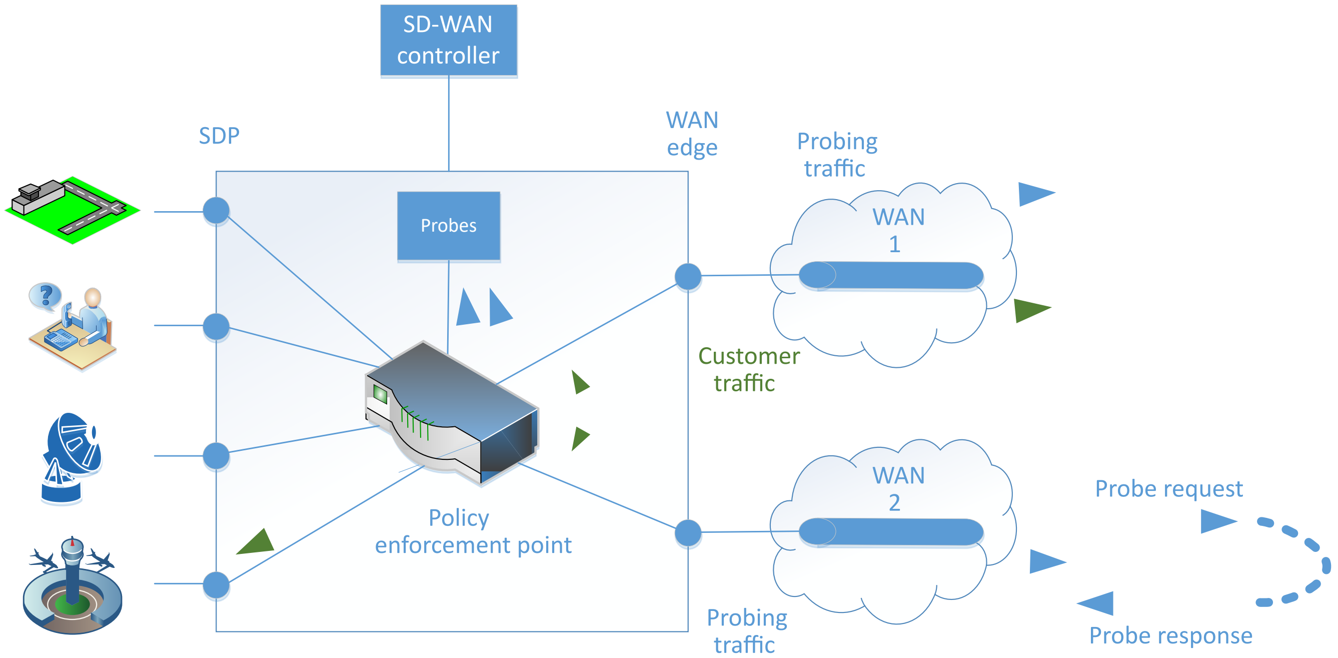

Approach 2: Multi-domain orchestration

An augmentation of the previous approach is the cross-layer orchestration, including the use case where the ANSP also operates (some of) the WANs. This scenario is generic, and may cover exotic combinations where the ANSP operates a terrestrial WAN, but relies on a telecommunication operator for a backup WAN. For the sake of simplicity, we consider the case where the ANSP operates everything. In this scenario, as can be observed in Figure 1b, two new components appear: (i) SDN controller(s) and (ii) orchestrator.

On the one hand, the SDN controllers are responsible for each of the WANs and optimise their performance without disturbing, if possible, the traffic from the overlay. Usually, this is of interest in cases where the underlay network provides simpler/faster mechanisms to detect and solve problems, than those possible with an SD-WAN-only solution, for example, with MPLS-TE and precomputed restoration paths may be possible to do sub-50-ms blackout-based steering, even before the overlay network is able to detect it.

On the other hand, the orchestrator plays the role of coordinator between the SD-WAN controller and the SDN controllers. The rationale is that having different mechanisms to do a similar task, for instance, to keep the end-to-end traffic performance as high as possible, may be counterproductive and promote even more issues.

Approach 3: Multilayer controller

The third approach is a simplification of the previous one, as all the domains (SD-WAN and SDN-controllers) are merged into one. The idea is that a single SDN controller is responsible for the whole end-to-end infrastructure. However, this approach is not even an offering in the enterprise telecommunication market, because of the operational complexity, and thus it is out of the scope of the paper. The reason for the lack of offering is that operators usually have different departments for each network segment and, maybe, technology. Therefore, it is not easy for solution providers to provide a unified platform.

Scalable and hierarchical SDN architectures

This section is intended to provide different deployment architectures for SD-WAN solutions, that is, approach 1 in the previous section. Approaches 2 and 3 are out of the scope of the paper.

First, several alternatives are presented to cover different use cases. Then, those solutions are compared in terms of different aspects such resilience, system overhead and agility to support new steering decisions.

As a reference scenario, let us assume a generic country-wide ANSP network with a set of flight information regions (FIRs) partitioning the national airspace. Each of these FIRs will be integrated by different locations such as area control centres (ACCs, or air route traffic control centres -ARTCCs- in US), approach/departure control (APP, or terminal radar approach control -TRACON- in US), remote sites including radar and air-ground (A/G) radio equipment, and so on.

For simplicity, we will assume that each site within a FIR, regardless of the centre/remote nature, is connected to, at least, a FIR-wide WAN. In case some remote location does not have capabilities to host a CPE like in Figure 2, we will assume that application data is aggregated in a CPE in another location so that, from the CPE perspective, it is a "local" application. Besides, in the event that a site has to be used to route traffic from another site, we will assume that the original application is also "local" here. For inter-FIR traffic, we assume a set of, this time, country-wide WANs interconnecting points of presence (PoPs) for each specific FIR (let us assume a single PoP per FIR for simplicity), where all WAN providers concur. Therefore, these last locations face intra-FIR WANs on the one side, and inter-FIR WANs on the other side, being thus gateways between FIRs.

All this information is represented in Figure 3, where FIR 2 and 3 are similar to FIR 1 but only PoPs are represented for simplicity purposes.

As can be seen in Figure 3, per-FIR partitioning may introduce another layer of complexity that is the inter-domain routing between FIRs. Consequently, inter-FIR traffic is omitted in the rest of the paper, and all traffic will be assumed to stay within the local FIR.

Moreover, the discussion about how many WANs are used for each site, or whether there is a redundant data plane (PEP) on each site, is out of scope. It could even be possible that a given application is attached to two different sites and it is able to select one itself, or support duplication/deduplication capabilities (i.e., linked session support in ED-137B[7]). This is a task that should be performed during the network design phase to determine the overall end-to-end availability for each and every application. Nevertheless, the SD-WAN controller may take care of monitoring and reporting service-level agreement (SLA) reports.

Figure 3: Country-wide ANSP network

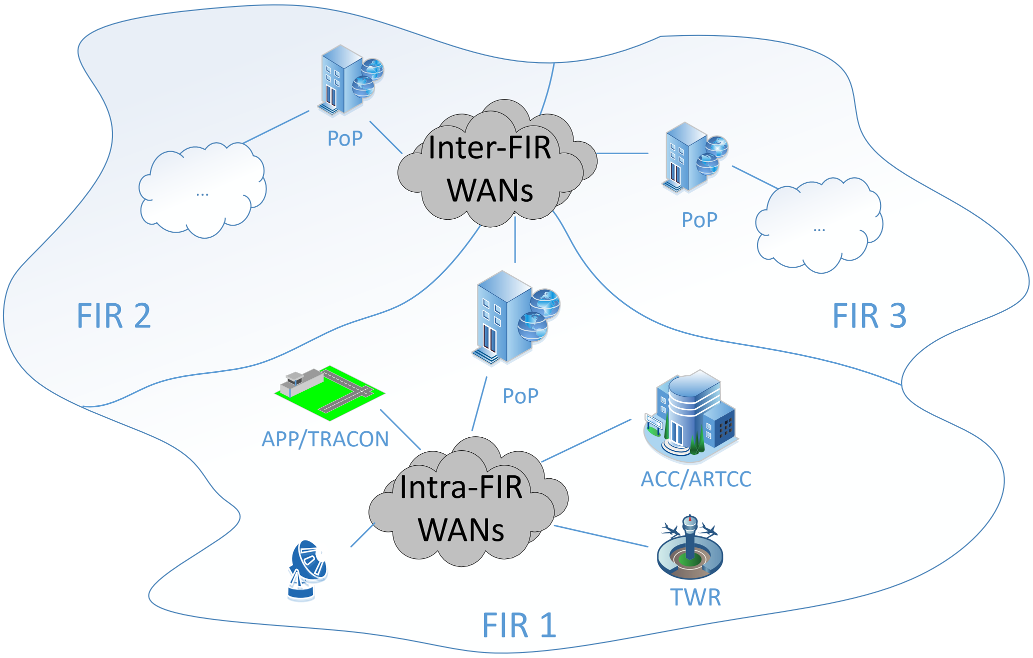

Approach 1: Per-site standalone

In this first proposal, there is one “isolated” SD-WAN controller sitting inside the CPE on each site, as another component. Essentially, each site takes care of its ingress traffic, that is, local traffic aiming to be forwarded to another site. There is no control plane mechanism exchange between sites, as each SD-WAN controller does not talk to any other in any site. As such, each site can only take local decisions, but can be the fastest approach since it does not have to coordinate with anybody else, as represented in Figure 2. By the use of proper probing capabilities (response to probing in one direction includes the request in the other direction), can reduce the influence of active probing on data traffic. This approach presents two main drawbacks:

- No mechanism to support symmetric traffic. It is not possible for bidirectional traffic (i.e., a voice conversation) to use the same WAN in both directions because the lack of communication between SD-WAN controllers in different sites.
- Single point of failure (SPOF) in the control plane in a per-site level. This means that if the SD-WAN controller fails, or there is a maintenance window, the data plane becomes unmanaged for some time, which may lead to degradation because of inability to steer traffic to a better path.

Given a flow F willing to communicate from site A to site B, the only task of the SD-WAN controller is to select which WAN to use for F in the A→B direction according to the probing results. If F is part of a bidirectional application (i.e., voice), the SD-WAN controller in site B takes its decision independently for the reverse F’ flow. This solution is illustrated in Figure 4.

Figure 4: Per-site standalone

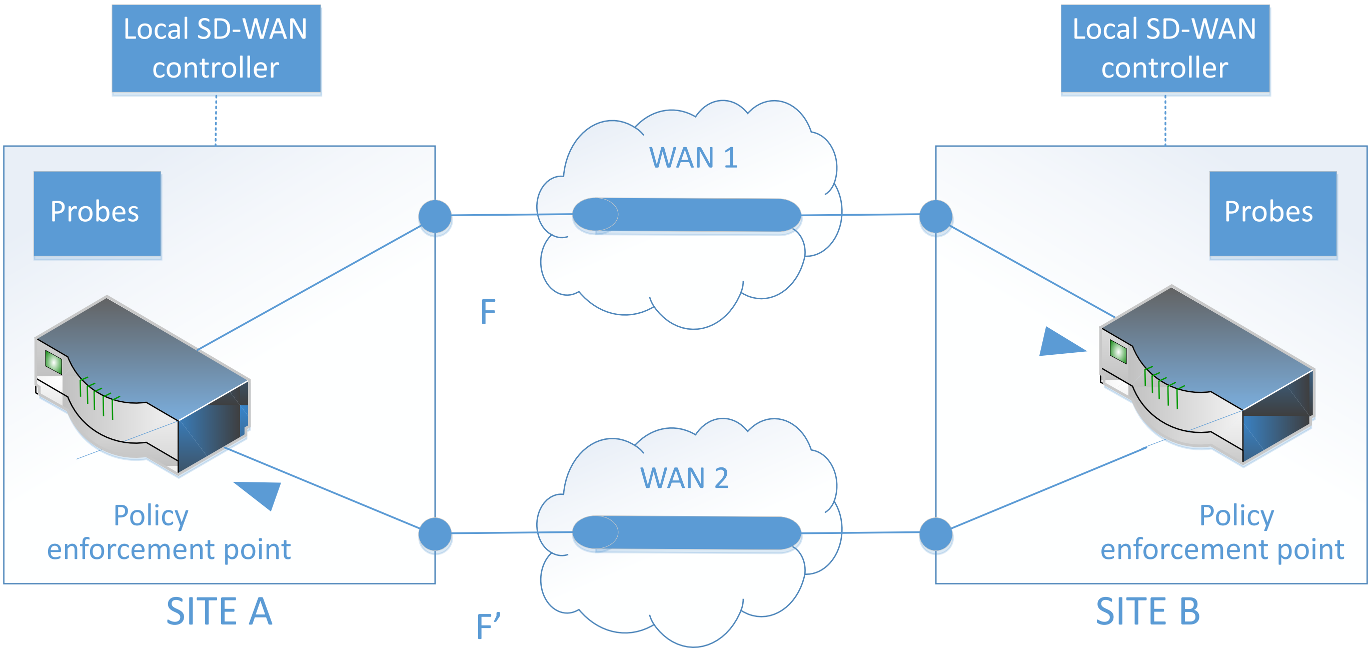

Approach 2: Per-site fail-over pair

To overcome the SPOF issue from the previous approach, this one assumes that there is a second SD-WAN controller per-site, implementing some sort of first-hop redundancy protocol and coordination mechanism to decide which one is used in an active-active setup. This capability is a key enabler for in-service system upgrade (ISSU), to maintain optimal WAN routing for the applications even during maintenance windows. Regarding flow routing the mechanism is the same as in the previous approach, with the additional requirement for both SD-WAN controllers to keep their internal datastores synchronised. This strategy is depicted in Figure 5.

Approach 3: Centralised (single)

The most prominent issue of the previous approach is the inability to perform symmetric routing, which is an important willing-to-have capability, if not mandatory, for voice communication systems in order to support proper dynamic delay compensation.

As a means to overcome this, a possible solution is to have a single centralised controller that coordinates all the operations from a single location. This approach has the advantage that the controller has a global view of whole the ATM overlay, and can perform symmetric routing decisions. However, there are two critical drawbacks making this solution impractical in a safety-critical environment:

- SPOF in a network-wide level. If the SD-WAN controller goes down, the whole overlay network becomes unmanaged.
- Detection and steering times are longer than in previous approaches, since probing data has to be delivered towards the central location and, then, the new instructions from the SD-WAN controller have to be communicated to the PEPs, which it may take essentially the same time.

Figure 5: Per-site fail-over pair

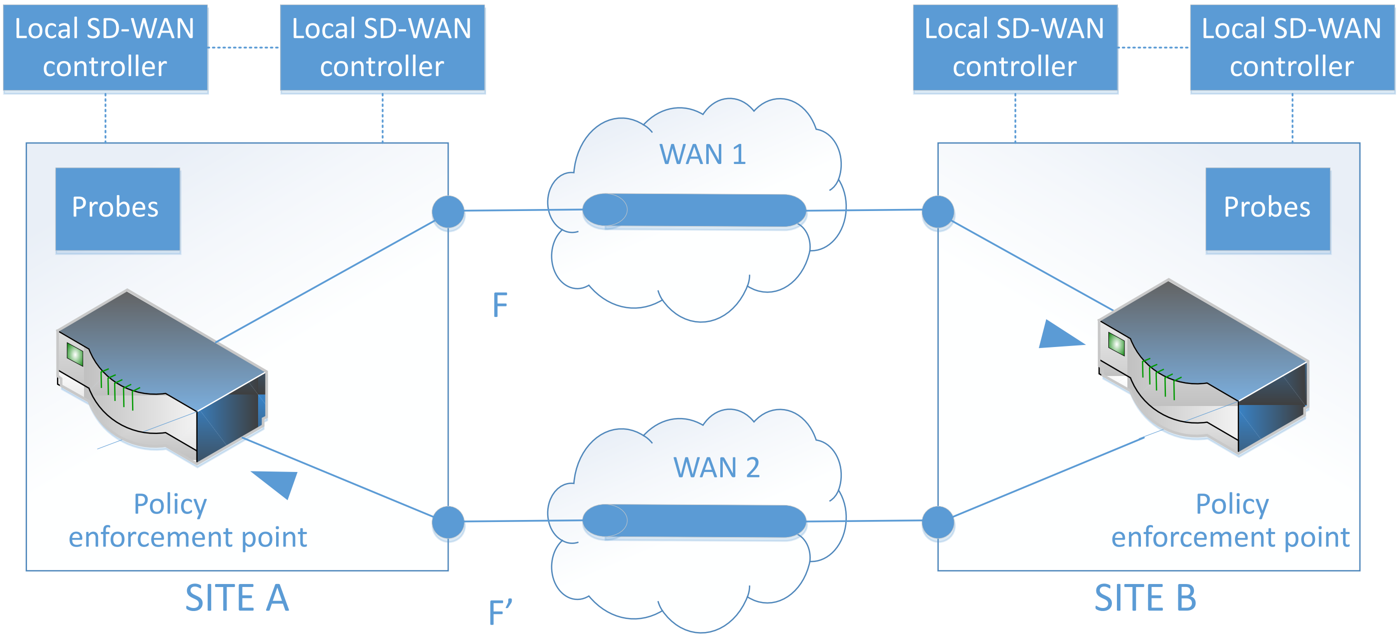

In this setting, a bidirectional flow F between sites A and B is steered by the SD-WAN controller according to the information collected from probes in sites A and B. Note that, if symmetric routing is required, the SD-WAN controller has to map QoS parameters for each candidate path in both directions, and then select the overall best, as shown in Figure 6. For unidirectional traffic, or without symmetricity requirements, there is no advantage, but the disadvantage of slower detection and steering times.

Approach 4: Centralised (cluster)

As an intermediate solution between a fully-decentralised approach and the single centralised approach, there is the possibility to have a centralised cluster.

In this setup, there is a cluster of SD-WAN controllers spread across each FIR. The difference between this approach and the standalone/fail-over pair one is that the controllers in the cluster setup have an (eventually consistent) synchronised datastore, but they can distribute the workload among the cluster to reduce detection and steering times.

The number of members in the SD-WAN controller cluster is usually limited to three, five, or seven, at most, being ideally lower than the number of sites, but always an odd number to avoid split-brain. Controller placement algorithms are used to determine the right number, and location, of cluster members.

In case of bidirectional traffic, both flows are assigned to the same cluster member to be able to make coordinated decisions. Besides, in case of failure or maintenance window, the workload is redistributed among the live cluster members.

Figure 6: Centralised (single)

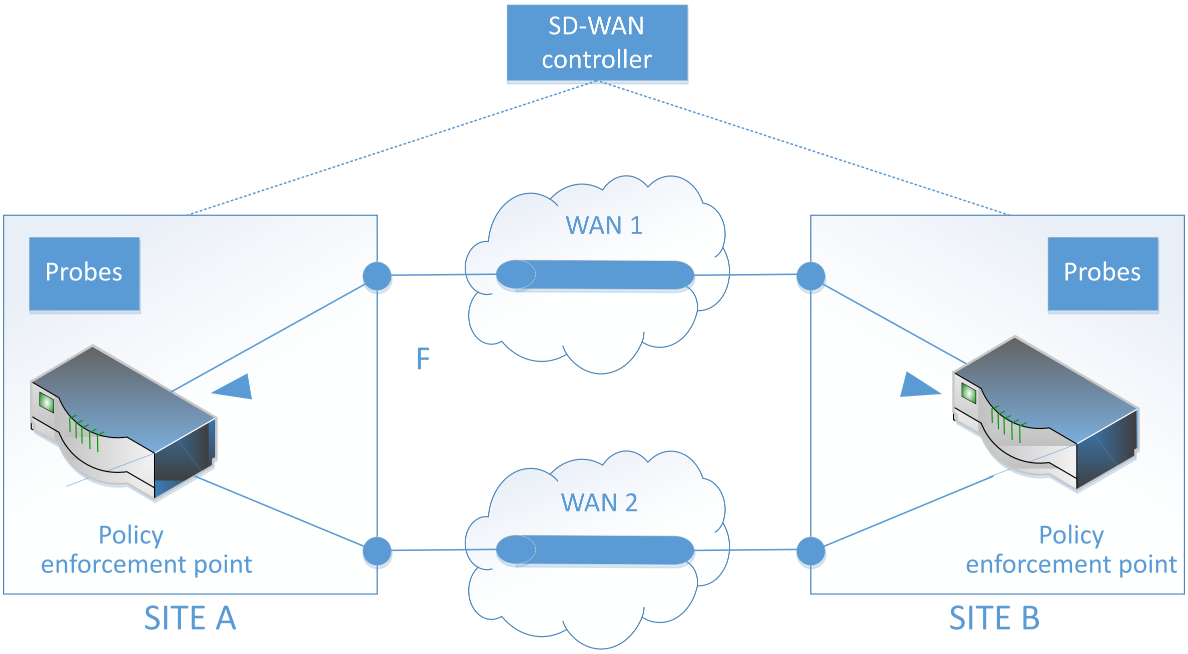

However, this approach still has the problem that if the network becomes partitioned, which means that not all sites are reachable from the set of live cluster members, some sites become unmanaged.

Compared to the previous scenario, in general detection and steering times are lower because it is expected that sites are assigned to the nearest cluster member. Whereas this statement may be challenged for bidirectional traffic.

In this setting, a unidirectional flow F between sites A and B is steered by the SD-WAN controller associated to site A according to the information collected from probes in site A. The only difference to the first two approaches is that the detection and steering time may be longer. By contrast, for a bidirectional flow F' between sites A and B, the steering decision is made by the cluster member associated to A←→B flows with probing results from sites A and B. This approach is presented in Figure 7.

Approach 5: Centralised with per-site fallback

The main outcome of the analysis of the previous approaches can be summarised in the following aspects.

On the one hand, for unidirectional traffic, that is, the one not requiring symmetric routing, the per-site approach provides the best performance in terms of detection and steering speed. On the other hand, for bidirectional traffic, that is, the one requiring symmetric routing, the centralised approach is mandatory in order to support coordination between both end sites.

The contribution of this new approach is the integration of a mechanism to overcome the partitioning issue of a purely centralised approach. Essentially, there is a two-layer hierarchy of SD-WAN controllers: (i) parent layer, covering the centralised approach; and (ii) child layer, covering the per-site approach. For the sake of simplicity in the explanation, we consider a single centralised controller in the parent layer, and a standalone controller per site in the child layer. Different proposals such as ACTN[8] (abstraction and control of traffic-engineered networks, proposed by the IETF) or Transport API[9] (defined by the Open Networking Foundation) aim to facilitate the implementation of this architecture.

Figure 7: Centralised (cluster)

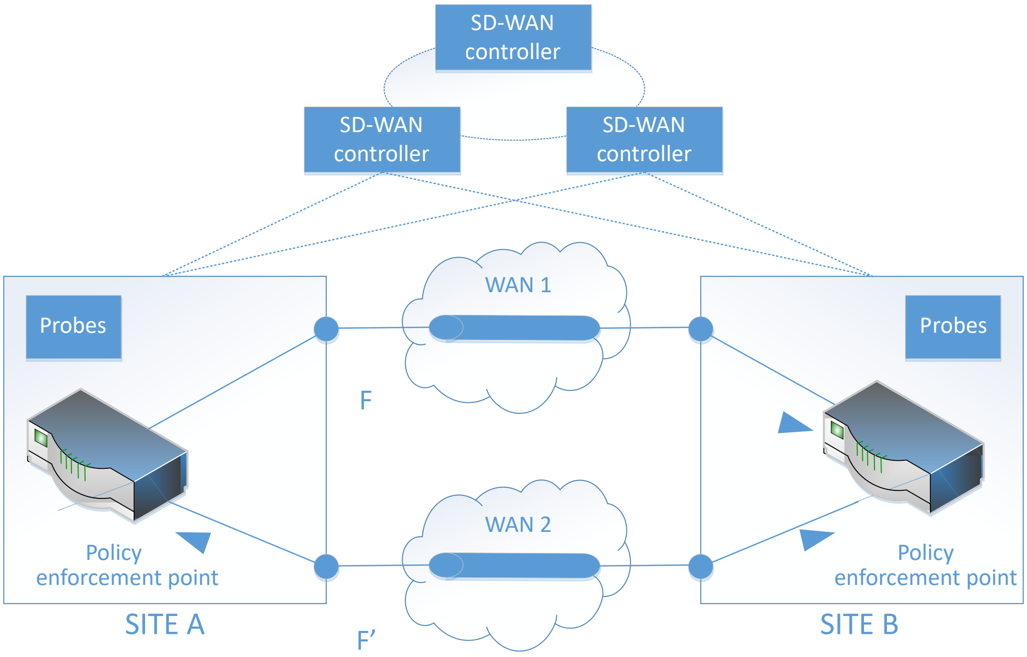

For this setup, we consider the parent layer responsible for traffic steering during normal operation. However, the parent controller does not have direct access to probes and PEPs, but to each individual local controller on each site. In this way, children controllers are a proxy between the parent controller and the data plane, being themselves passive observers. Therefore, both unidirectional and bidirectional flows are properly managed by the system, as shown in Figure 8.

Conversely, in the event that the control plane is partitioned, and connection between parent and children is lost, the local controller takes over its local (ingress) applications until the connection to the parent controller is restored. During the failure time, bidirectional traffic is not supported anymore, and such flows are split into two separate flows, each of them routed individually, as depicted in Figure 9.

Figure 8: Centralised with per-site fallback (failure-free scenario)

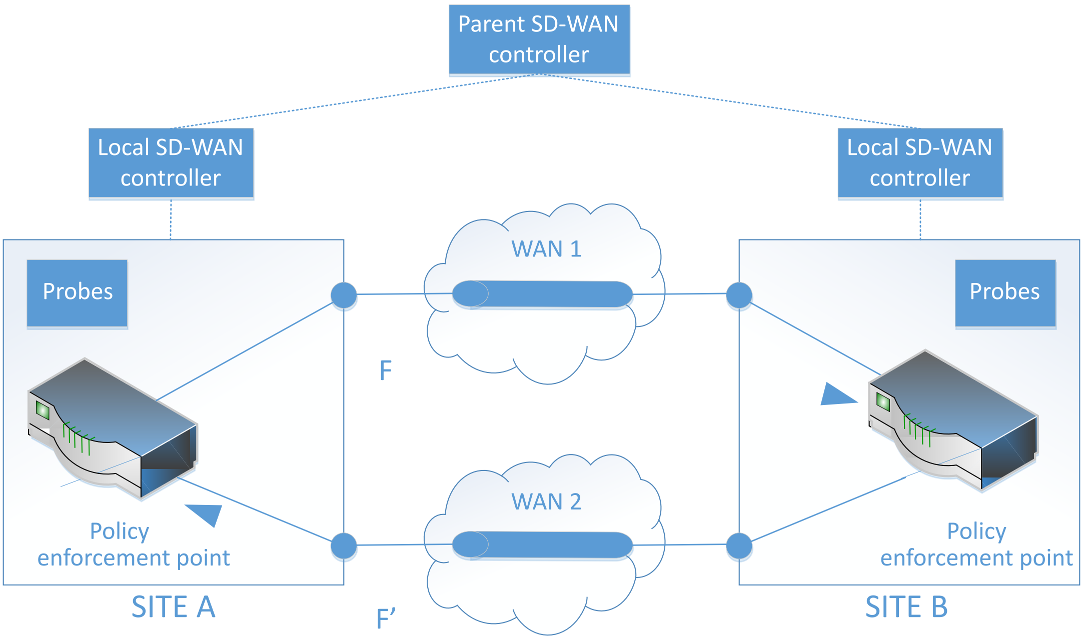

Figure 9: Centralised with per-site fallback (failure scenario)

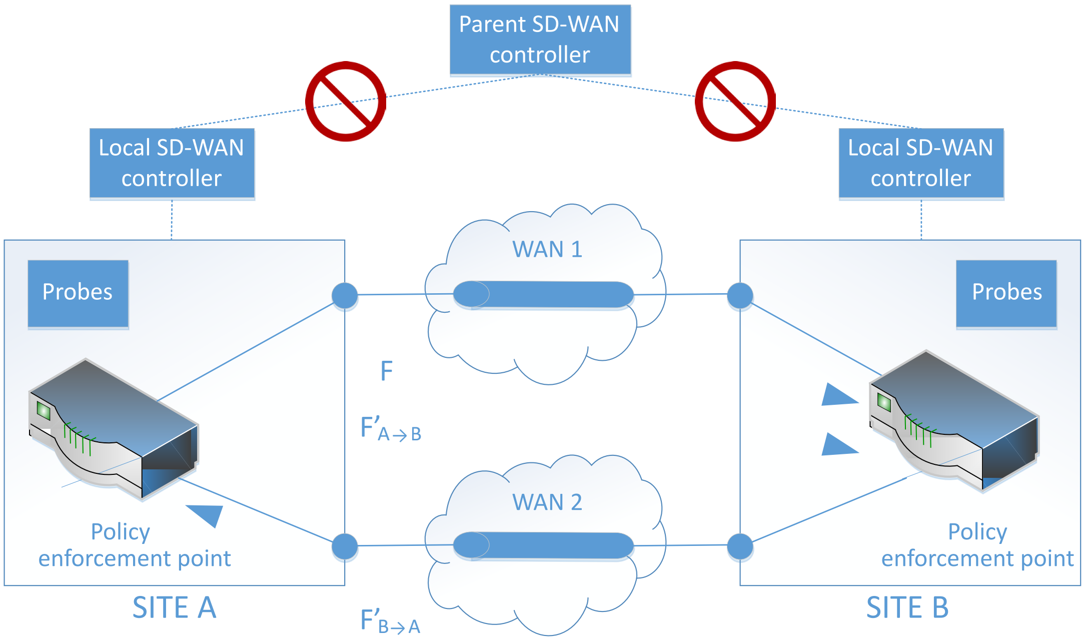

It is important to highlight that this configuration is expected to support some "delegation" functionality, as well. The reason is that the control plane may become partitioned partially, for example, only site A from Figure 9. In this situation, the expected behavior is that the parent controller delegates all the flows towards site A to the corresponding other sites, since it does not make sense to keep them centralised if one of the two members of the party is not reachable.

This delegation feature can be also applicable to sites with unidirectional-only traffic, since the centralised approach introduces extra overhead that is not needed: (i) probing results do not have to be delivered to the centralised controller; and (ii) detection and steering times are shorter because of local decisions. Note that mixing in one site flows handled by the parent controller with some others handled by the local controller may be tricky, as they can interfere and promote flapping events.

Supporting ATN/IPS

The ATN/IPS profile identifies a minimum set of requests for comments (RFCs) from the IETF that must be implemented by all airborne and ground nodes supporting aeronautical safety communications. This profile covers the transport and network protocol layer functions, including security, naming, address discovery and resolution, routing, mobility, multi-link and network management.

In the ATN/IPS architecture the airborne end-systems hosted on an aircraft are part of an IPv6 network connected to the ground network by one or more airborne routers (A-R)[10]. A-Rs have multiple radio interfaces that connects them via various radios infrastructures (e.g., SATCOM, LDACS, AEROMACS) to a given radio region on the ground. Typically, an A-R has a corresponding ground-based access router (AC-R) that terminates the radio protocol with the A-R and provides access services to the ground-based portion of the radio network infrastructure. Each radio region is interconnected with the ATN/IPS ground network via an air-to-ground router (A/G-R), which routes the traffic from/to ground end-systems via ground-to-ground routers (G/G-Rs).

The ATN/IPS ground network infrastructure is the internetworking region located between the A/G routers and the G/G routers.

Frequentis is currently validating the different aspects of ATN/IPS with a special focus on mobility and multilink as part of the SESAR Future Communication Infrastructure (FCI)[11] [12] [13]. The main goal is to avoid any impact of mobility on the airborne equipment by solving it in the ground network infrastructure. Therein, ground-based LISP (GB-LISP) has been chosen as the most promising solution for mobility in the ground infrastructure. Validations are performed from April to June 2019, and will be made available for the next edition of ICNS in 2020.

LISP is a routing solution that supports multi-homing and mobility, based on the semantic separation of IP address into Endpoint IDentifiers (EIDs) and Routing LOCators (RLOCs)[14]. RLOCs are used by LISP routers to create routing tunnels, whereas EIDs are used to identify non-publicly routable end devices. In this way, end-customer LISP functionality is deployed exclusively on customer endpoint routers, which perform both the egress tunnel router (ETR) and ingress tunnel router (ITR) functions of a LISP device (abbreviated as xTR).

This protocol enables xTRs to exchange mapping information between EIDs and RLOCs, analogous to the DNS for the resolution of IP addresses. Additionally, it defines directives to encapsulate/decapsulate IP packets toward EIDs using RLOCs as tunnel endpoints, seamlessly operating with the current TCP/IP stack. Nonetheless, LISP data packets are UDP packets whose payload contain original packets, whereas EIDs and RLOCs are syntactically equal to IPv4/IPv6 addresses.

Due to the separation between topology (RLOC) and identifier (EID), two control plane services are required: (i) map register (MR); and (ii) map server (MS). The former is the one receiving updates from RLOCs informing about "local" EIDs. The latter is the one answering queries from RLOCs about the RLOC to reach an EID.

In the GB-LISP architecture, a LISP overlay is layered over the ATN/IPS internetworking region (that is in the LISP RLOC space) and provides connectivity between end systems (that are in the LISP EID space) hosted in the aircrafts and those ones in the ground. The A/G-Rs and the G/G-Rs assume the role of LISP xTRs supported by a LISP mapping system infrastructure.

By considering the centralised architecture with per-site fallback, we can envision a DNS-like architecture in which there is a hierarchy of master-slave MS/MRs, each embedded into each of the SD-WAN controllers deployed in the CPEs and central locations. The overall architecture is illustrated in Figure 10.

Conclusion

In this paper, we have discussed different approaches for the introduction of the software-defined networking (SDN) paradigm into the air traffic management (ATM) world. Leveraging the concept of software-defined wide-area networks (SD-WAN), it is possible to create a solution able to cover essential requirements of a converged, safety-critical network, such as application-based traffic flow management based on differentiated QoS profiles and assessment via active probing, all based on an SDN overlay built on top of existing, heterogeneous enterprise WAN offerings.

By analysing different SD-WAN deployment options, we concluded that a hierarchical approach based on a centralised, parent controller layer with local fallback based on per-site child controllers is a resilient approach to avoid unmanaged network partitions under failure situations.

Further studies will cover new concepts like parent-child delegation for unidirectional flows capabilities, whereas the parent is used for bidirectional flows due to the global view in failure-free scenarios, or multi-domain SD-WAN deployments, including benchmarking on a real testbed. Finally, we have briefly described the binding between the SD-WAN solution and ground-based LISP, as part of the Future Communications Infrastructure (FCI) initiative of the SESAR 2020 programme.

Acknowledgements

The work related to ATN/IPS integration is based on the Project 14.2.4 of the SESAR2020 Joint Undertaking under the European Union's Horizon 2020 Research and Innovation programme.

Authors

Jose-Luis Izquierdo-Zaragoza, Wolfgang Lins, Peter Leydold, François d'Humières, Frequentis AG, Vienna, Austria
Dieter Eier, Frequentis USA, Columbia, MD

Figure 10: GB-LISP-enabled SD-WAN solution

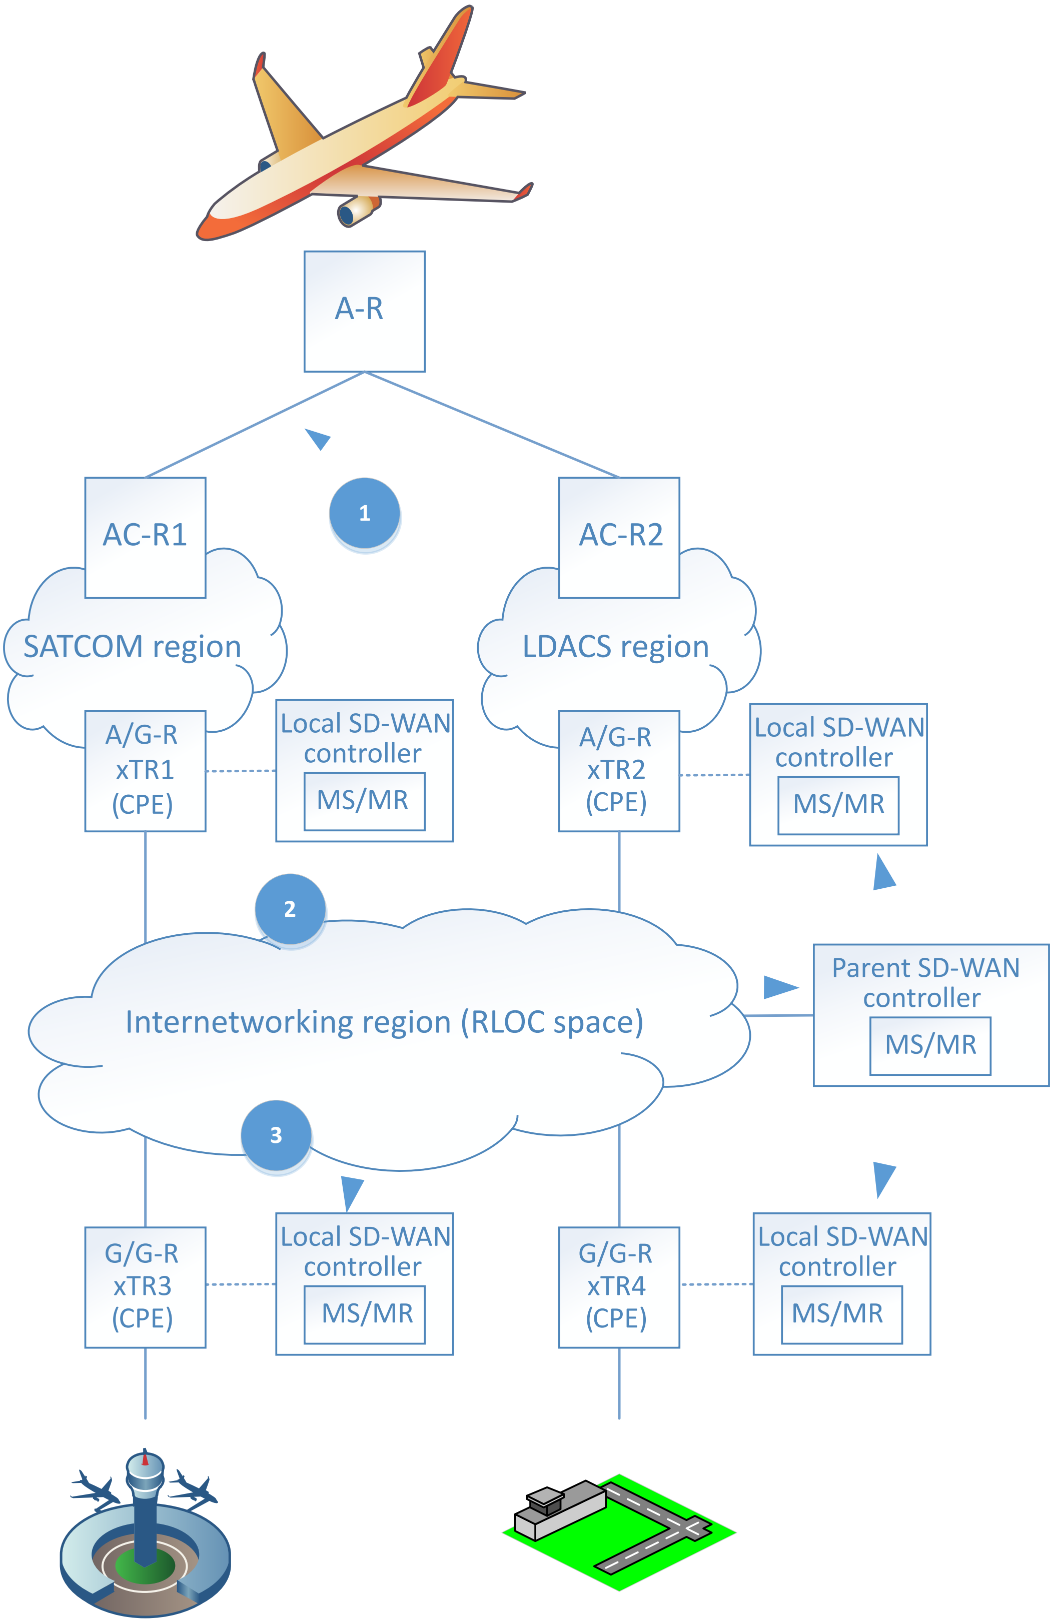

Endnotes

1. EUROCONTROL. 2013. Determining the benefit and costs of Centralised Services. Skyway, no. 60 (Winter 2013)", pp. 40-41.

2. RTCA. 2011. DO-278A – Software Integrity Assurance Considerations for Communication, Navigation, Surveillance and Air Traffic Management (CNS/ATM) Systems.

3. EUROCAE. 2012. ED-109A – Software Integrity Assurance Considerations for Communication and Navigation and Surveillance and Air Traffic Management (CNS/ATM) Systems.

4. EUROCAE. 2009. ED-153 – Guidelines for ANS Software Assurance.

5. ICAO. 2015. Doc. 9896 – Manual on the Aeronautical Telecommunication Network (ATN) using Internet Protocol Suite (IPS) Standards and Protocol. 2nd edition.

6. Katz, Dave, Dave Ward. 2010. Bidirectional Forwarding Detection (BFD). RFC 5880.

7. EUROCAE. 2012. ED-137B – Interoperability Standards for VoIP ATM Components.

8. Ceccarelli, Daniel, Young Lee. 2018. Framework for Abstraction and Control of Traffic Engineered Networks. IETF draft draft-ceccarelli-teas-actn-framework-15.

9. Open Networking Foundation. TR-527 – Functional Requirements for Transport API.

10. Haindl, Bernhard, Manfred Lindner. 2016. Ground Based Lisp for Multilink Operation in ATN/IPS Communication Infrastructure. IEEE/AIAA 35th Digital Avionics Systems Conference (DASC 2016). Sacramento (CA), United States.

11. SESAR. Future Communications Infrastructure (FCI) Initial Concept Description. SESAR2020 PJ14.02.04 Deliverable D5.1.020.1 (Ed. 00.01.02).

12. SESAR. Future Communications Infrastructure (FCI) Initial Transversal and Complementary Studies. SESAR2020 PJ14.02.04 Deliverable D5.1.020.2 (Ed. 00.01.03).

13. SESAR. Future Communications Infrastructure (FCI) Functional Requirements Document. SESAR2020 PJ14.02.04 Deliverable D5.2.010 (Ed. 00.01.02).

14. Farinacci, Dino, Vince Fuller, Dave Meyer, Darrel Lewis. 2013. The Locator/ID Separation Protocol (LISP). RFC 6830.

FREQUENTIS AG
Innovationsstraße 1
1100 Vienna, Austria
Tel: +43-1-811 50-0
www.frequentis.com

The information contained in this publication is for general information purposes only. The technical specifications and requirements are correct at the time of publication. Frequentis accepts no liability for any error or omission. Typing and printing errors reserved. The information in this publication may not be used without the express written permission of the copyright holder.