

On Building Secure Wide-Area Networks over Public Internet Service Providers

Marc Wyss

ETH Zurich
Department of Computer Science
Zurich, Switzerland
marc.wyss@inf.ethz.ch

Roland Meier

armasuisse Science and Technology
Cyber-Defence Campus
Thun, Switzerland
roland.meier@ar.admin.ch

Llorenç Romá

armasuisse Science and Technology
Cyber-Defence Campus
Thun, Switzerland
llorenç.roma@ar.admin.ch

Cyrill Krähenbühl

ETH Zurich
Department of Computer Science
Zurich, Switzerland
cyrill.kraehenbuehl@inf.ethz.ch

Adrian Perrig

ETH Zurich
Department of Computer Science
Zurich, Switzerland
adrian.perrig@inf.ethz.ch

Vincent Lenders

armasuisse Science and Technology
Cyber-Defence Campus
Thun, Switzerland
vincent.lenders@ar.admin.ch

Abstract: Many public and private organizations use wide-area networks (WANs) to connect their geographically distributed sites. Given that these WANs are often critical for the organization's operations, their security with respect to confidentiality, integrity, and availability is crucial.

A high level of security can be reached if the WAN is built with a dedicated network infrastructure, with the organization operating its own layer-2/3 routing, for example, multiprotocol label switching on top of dedicated fibers or leased lines. Unfortunately, this approach is often slow to deploy, requires high operational effort, and is too expensive for many use cases.

A cheaper alternative is to construct the WAN as an overlay network on the infrastructure of public Internet service providers (ISPs), for example, using virtual

private network tunnels between the sites. Unfortunately, the security of such a WAN is suboptimal. For instance, traffic analysis attacks (on encrypted traffic) can reveal sensitive information transmitted over these public networks, compromised routers between the sites can alter packets, and network-layer distributed denial-of-service (DDoS) attacks can disrupt connectivity.

In this paper, we explore a novel inter-ISP network architecture that provides the desired level of control and security for WAN operators, achieving the best of the two above approaches: strong security properties on a cost-efficient public Internet fabric. Our architecture builds on the SCION next-generation Internet architecture and adds extensions for fine-grained path control, connectivity guarantees in the presence of DDoS attacks, and traffic analysis prevention. With this architecture, WAN operators can build on public layer-3 network connectivity services to deploy secure WANs.

Keywords: *wide-area networks (WANs), WAN security, SCION*

1. INTRODUCTION

Wide-area networks (WANs) are often used to connect multiple sites of a large organization. To protect data sent between sites, the WAN must ensure the confidentiality, integrity, and availability (called the CIA triad) of inter-site communication. Ideally, a WAN is built on a dedicated infrastructure consisting of trustworthy network devices and is fully under the control of the organization itself. In practice, WAN deployments frequently make use of leased lines to ensure traffic isolation and to provide quality of service (QoS) guarantees. Unfortunately, building such a WAN requires all sites to be connected via leased lines, which may be prohibitively expensive.

As an alternative, sites can connect to public Internet service providers (ISPs) and use overlay connections over the public Internet which are protected using technologies for virtual private networks (VPNs), such as IPsec tunnels. Such an ISP-based deployment significantly reduces the cost as (1) ISP connectivity typically has a lower cost than leased-line connectivity, and (2) every site only requires a single access point at its upstream ISP. Unfortunately, Internet overlay connections have much weaker security properties compared to leased lines. Thus, they are typically insufficient for the requirements of organizations requiring high security, such as critical infrastructure providers, governmental bodies, and military organizations. For example, data confidentiality is hindered by eavesdropping and traffic hijacking, data integrity may not be preserved due to the presence of adversarial network devices that


inject or modify packets, and communication availability cannot be guaranteed due to distributed denial-of-service (DDoS) attacks. Furthermore, since the infrastructure is shared, heavy load from other customers can impair the availability of the WAN connections due to network congestion. Centralized solutions, such as software-defined WAN (SD-WAN) [1] approaches and DDoS mitigations based on rerouting and scrubbing [2], can address some of these issues. However, these solutions require a single entity with direct control over the complete network infrastructure, which hampers deployment in widely distributed networks, such as those spanning multiple countries. Furthermore, in a federated setting, finding a single entity trusted by all actors is not always possible.

In this paper, we analyze the challenges of building a secure WAN and propose a WAN architecture that achieves security properties similar to a leased-line deployment—but only requires deploying a single ISP connection per site, resulting in a significant cost reduction. These properties can be achieved by leveraging a recent trend called path-aware networking, which allows endpoints to gain insight into the forwarding paths taken by their traffic (path transparency) and influence the forwarding path choice according to their criteria (path control). Figure 1 compares a leased-line deployment with our proposed secure WAN deployment from the viewpoint of a specific site. In particular, we make use of the next-generation Internet architecture SCION [3], which provides path transparency and control at the level of autonomous systems (ASs), representing individual network operators. Furthermore, we leverage various recent network protocols to achieve the CIA triad, for example, using bandwidth reservations for strict QoS guarantees, DDoS protection, and traffic obfuscation.

Contributions: Our paper makes the following contributions:

- We analyze the security challenges of building a secure WAN over layer-3 connectivity services from ISPs and derive the security properties that an inter-ISP network architecture must support (Section 2).
- We propose an architecture that combines new network protocols to achieve the required security properties (Section 3).
- We present our initial deployment comprising four WAN sites in two countries connected over four ISPs and discuss a potential use case for multinational collaboration (Section 4).

FIGURE 1: DEPLOYMENT SCENARIOS FOR SECURELY CONNECTING MULTIPLE SITES FROM THE PERSPECTIVE OF SITE 3 (CONNECTING TO ALL OTHER SITES). IN THE FIRST SCENARIO, THE SITES ARE CONNECTED VIA PAIRWISE LEASED-LINE CONNECTIONS, AND IN THE SECOND SCENARIO VIA OUR PROPOSED SECURE WAN SOLUTION. TRAFFIC BETWEEN SITE 3 AND SITE 4 SHOULD NOT LEAVE A CERTAIN JURISDICTION (GREEN AREA)—FOR EXAMPLE, TO PRESERVE POLICY COMPLIANCE—AND ALL TRAFFIC SHOULD AVOID A CERTAIN UNTRUSTED NETWORK OPERATOR (RED)



2. SECURITY CHALLENGES

Our aim is to craft a secure WAN solution tailored to the public Internet. In this context, “secure” encompasses the CIA triad—confidentiality, integrity, and availability—which is crucial for WANs: Confidentiality in the WAN context means safeguarding all traffic, both metadata and payload, and minimizing data proliferation. Integrity involves ensuring that data remains unaltered during transmission and enabling the detection of unauthorized tampering. Availability ensures seamless communication, allowing successful transmission of all traffic up to a specific bandwidth threshold between all WAN sites.

Table I highlights significant threats against those security properties and presents possible mitigation measures. We do not claim that our identification of threats is complete, but the threats discussed here represent our best estimate of the most significant challenges.

TABLE I: THREATS AGAINST THE CONFIDENTIALITY, INTEGRITY, AND AVAILABILITY OF WAN COMMUNICATION AND POSSIBLE MITIGATION MEASURES. SOME MITIGATION MEASURES CAN HAVE THE SIDE EFFECTS OF CONTRIBUTING TO MITIGATING AND/OR AMPLIFYING CERTAIN THREATS

Mitigations → Threats ↓		Traffic encryption	Traffic shaping and padding	Traffic filtering	Traffic authentication	Path authentication	Path control	Traffic prioritization
Confidentiality	Eavesdropping (payloads)	V					V/X	
	Eavesdropping (metadata)	V	V				V/X	
	Traffic hijacking					V	V	
Integrity	Traffic injection				V			
	Traffic modification				V			
Availability	Traffic dropping			X			V	
	Traffic hijacking					V	V	
	Congestion		X				V/X	V
	Volumetric DDoS			V	V		V	V
	Topology changes						V	

- V Mitigates the threat
- V Can contribute to mitigating the threat
- X Can contribute to amplifying the threat in general
- X Can contribute to amplifying the threat in general but does not apply to the architecture proposed in Section 3

A. Threats

Let us now discuss the threats listed in Table I in more detail.

1) Confidentiality

Confidentiality on the Internet covers both the secrecy of communication and the requirement of network operators to hide sensitive parts of their network topologies. *Eavesdropping*, for example, involves unauthorized interception of data in the network, potentially compromising confidentiality by allowing malicious entities to access sensitive information without the knowledge of the sender or receiver. If traffic is encrypted, an adversary cannot directly infer sensitive information from the packets' payload, but even eavesdropping on encrypted traffic can be problematic, as it can reveal metadata that enables traffic analysis.

Traffic analysis concerns the analysis of patterns, behaviors, or characteristics within data transmission, posing a risk to confidentiality by allowing the inference of sensitive information, even if traffic is encrypted and source or destination addresses are unknown. Traffic analysis can, for example, infer visited websites, streamed videos, or voice over IP (VoIP) calls from traffic volume, packet sizes, or timing information [4]. Furthermore, as an emerging technology, quantum computing presents a looming threat to current cryptographic algorithms. It has the potential to compromise the confidentiality of previously recorded encrypted traffic by breaking encryption algorithms that are currently considered secure.

Traffic hijacking allows an off-path adversary to redirect traffic across its network nodes, to essentially become an on-path adversary and gain access to confidential communication flows. An adversary can then eavesdrop and perform traffic analysis. This enables powerful eavesdropping attacks as it significantly increases the attack surface.

2) Integrity

Traffic injection, commonly referred to as packet spoofing, involves creating and transmitting network packets with falsified source addresses or other misleading information. It jeopardizes integrity by potentially circumventing defense systems, allowing for the exploitation of vulnerabilities, and enabling many different types of denial-of-service attacks. Additionally, control messages such as from the Internet control message protocol can be spoofed, leading to misinformation and potentially disrupting network operations.

Traffic modification refers to the unauthorized alteration of data packets during transmission, directly compromising integrity. Packets may be modified due to various factors such as transmission errors, malicious manipulation, or faults within packet processing systems. A common attack that leverages traffic modification is a man-in-the-middle attack, where an adversary places itself between the communicating endpoints and modifies selected packets. Quantum computing may impact traffic integrity by potentially breaking asymmetric cryptography used in digital signatures.

3) Availability

Achieving availability is often considered more challenging than confidentiality or integrity [5].

Traffic drop, such as where an on-path attacker deliberately drops either selected or simply all packets, is one of the most difficult threats to mitigate.

Path hijacking involves malicious entities diverting network traffic away from its intended path. It impacts availability by rerouting traffic through unauthorized paths, leading to potential delays or packet loss. Off-path attackers often use path hijacking attacks to essentially become on-path attackers. Path hijacking attacks are particularly powerful in combination with traffic-dropping attacks.

Congestion, whether naturally occurring due to high network demand or because of volumetric DDoS attacks, can affect both network and server availability. Network congestion leads to delays, packet loss, and decreased throughput, hindering communication, while DDoS attacks against servers can render their hosted services inaccessible to legitimate users. Volumetric DDoS attacks are a major reason why WAN operators prefer using leased lines to communicating over the public Internet, a trend amplified by the fact that attacks have grown in strength over recent years [6], [7]. Well-connected entities with links at hundreds of gigabits of capacity, as well as distributed botnets, sometimes leveraging powerful cloud virtual machines, can execute such attacks at a large scale [8]. Additionally, the proliferation of 10 Gbps home connections has made volumetric DDoS attacks easier and even more potent.

Topology changes regarding the network's physical or logical structure can potentially disrupt established communication paths. Such changes not only comprise topology modifications such as rewiring but also link or router failures. Link failures, whether due to hardware issues or physical damage, compromise all communications that use the link as part of their forwarding path. Router failures, whether caused by power outages, hardware malfunctions, natural disasters, or malfunctions or bugs in networking software, similarly affect availability. In such cases, data flows might be interrupted or rerouted; however, routing protocols may take time to converge and find the best available paths, as routers need to adjust and reach a consistent view of the network. During this convergence period, there can be temporary inconsistencies in routing information resulting in poor QoS, including high latency or jitter, and even lost connectivity, rendering services communicating over the WAN unusable.

B. Mitigation Measures

In this section, we discuss some possible mitigation measures for threats to a secure WAN architecture.

1) Traffic Encryption

Implementing encryption—for example, using transport layer security or IPsec—ensures that data transmitted across the network remains illegible to unauthorized entities. Also, employing multiple independent layers of encryption can effectively diminish the risk associated with misconfigurations and vulnerabilities. In light of the

impending threat posed by quantum computing, the adoption of post-quantum secure encryption becomes paramount.

2) Traffic Shaping and Padding

Traffic shaping involves regulating the flow of data toward a consistent and controlled transmission rate. Smoothing out data bursts or adding chaff packets can make traffic patterns less detectable. Thus, it becomes harder for an eavesdropper to discern specific patterns or extract meaningful information from a shaped traffic flow. In particular, shaping obscures the timing and volume of data transmission, thereby making it more challenging for adversaries to determine the nature and content of the communication. Padding involves adding extra data to packets to obscure the actual size or structure of the transmitted packets, rendering eavesdropping less effective. While shaping may delay the delivery of data, padding introduces additional data that occupies network bandwidth without conveying useful information and thus has the side effect of wasting resources.

3) Traffic Filtering

Network devices apply traffic filtering on different layers. One example is a firewall, which may filter packets at the transport and network layer or perform deep packet inspection to check application-layer data. By creating a boundary between a network and the Internet, a firewall can effectively protect against external network scanning while still allowing the network operator to debug the network from inside. Furthermore, firewalls can protect services hosted within the network from outside adversaries by rate-limiting incoming requests. Traffic filtering is also used for defensive mechanisms at end hosts that safeguard against DDoS attacks, ensuring that benign traffic can always successfully reach the destination application without being dropped due to congestion at the receiving end. Depending on the filtering technique, traffic filtering can have a negative impact on network availability if traffic is dropped by mistake (in false positive classifications).

4) Traffic and Path Authentication

The authentication of traffic is important both in the control plane and in the data plane.

Authentication in the control plane ensures that adversaries cannot tamper with control plane messages, such as route announcements sent by ASs and address prefix authorization messages issued by a public key infrastructure (PKI), thus preventing path hijacking attacks. Additionally, if control messages sent to endpoints are authenticated, an adversary cannot produce fake control messages, which may be used to revoke existing paths or inject non-existing devices in a traceroute reply.

Finally, by authenticating control messages originating at endpoints, a firewall can reply to control messages from authorized endpoints only, allowing networks to hide topology information.

In the data plane, authentication is used to ensure the integrity of end-to-end traffic, where authentic keys are typically fetched from a trusted PKI. Source authentication also prevents many different types of DDoS attacks that rely on spoofed source addresses, for example, reflection and amplification attacks, because firewalls can drop unauthenticated traffic.

5) Path Control

Path control, requiring a certain level of path transparency, ensures that traffic follows predefined routes, preventing it from leaving designated areas or regions. From an organization's WAN perspective, fine-grained path transparency and control are desirable because this enables traffic steering through trusted routers. Furthermore, path control mitigates path hijacking, as an attacker cannot influence the forwarding path through routing attacks. It also enables routing traffic around failing or congested links and routers, and around infrastructure believed to eavesdrop and perform traffic analysis. In addition to selecting preferred paths, the network should also provide path validation, that is, a mechanism to ensure that traffic is indeed sent over the selected path. Path validation can further be strengthened by remote attestation of on-path routers to ensure that they exist, are correctly configured, and run the intended software. Finally, path control can mitigate the threat of quantum computing by circumventing an adversary such that it cannot record and later decrypt the data.

In addition to selecting a desirable path to send traffic, endpoints can also leverage path control to enable resource pooling and simultaneously send data across multiple paths. This allows endpoints to achieve higher throughput but comes with its own challenges. The sender must ensure that no adversary is located on any of the selected paths to protect against eavesdropping and that other non-multipath flows do not experience excessive congestion, for example, using fair multipath congestion control protocols [9].

6) Traffic Prioritization

Traffic prioritization improves the QoS of priority traffic by instructing on-path routers to forward it with a higher priority compared to other traffic. This allows network operators to support low-latency and low-jitter communication while simultaneously maximizing the overall link usage for latency-insensitive best-effort traffic, such as file transfers.

A specific application of traffic prioritization is to explicitly reserve the required bandwidth for the priority traffic. Allocating bandwidth along the forwarding path between two WAN sites using an inter-domain bandwidth reservation protocol [10] can ensure that critical data flows smoothly through the network, even during periods of high demand or congestion. As traffic sent over a bandwidth reservation is unaffected by congestion, it shows properties similar to leased lines.

3. PROPOSED ARCHITECTURE FOR SECURE WANS

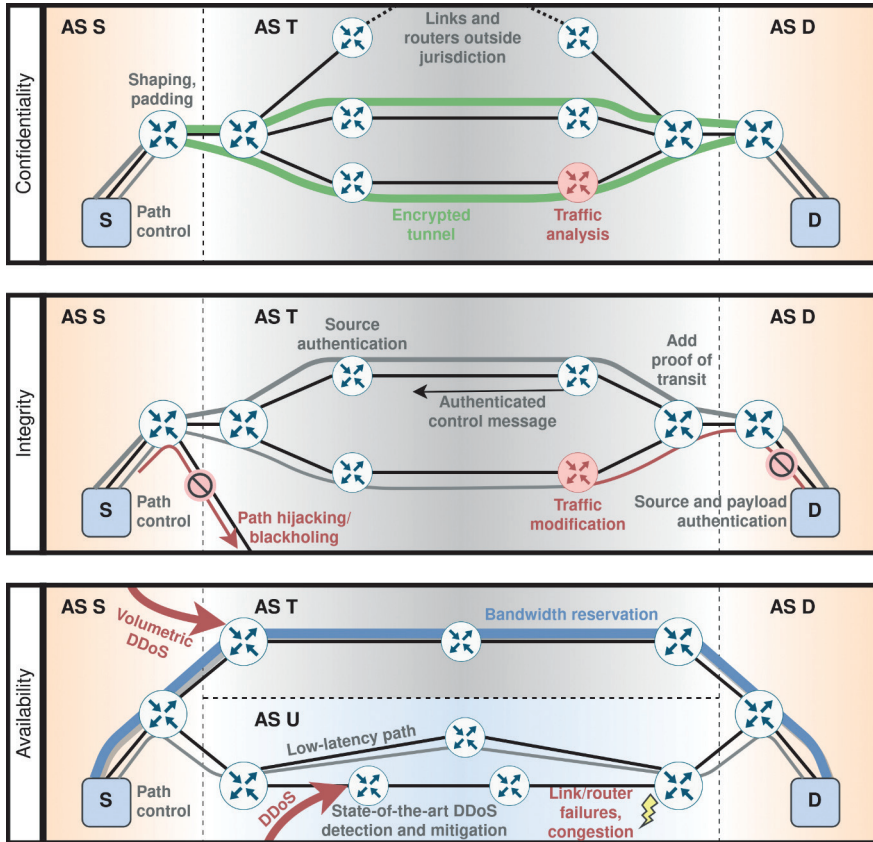
Considering the threats and the methods of mitigating them outlined earlier, we present an architecture that integrates a concise selection of mitigation measures optimized for WAN operators. Our goal is to streamline deployment complexity at ISPs by minimizing the number of mitigation measures while ensuring comprehensive coverage against the identified threats. While the selected mitigation measures address a wide range of threats, our architecture remains flexible, allowing for the inclusion of supplementary measures as needed.

A. Overview

The core component of our architecture is the next-generation Internet architecture SCION [3], which provides AS-level path control and authenticated control messages. SCION serves as the foundation for several systems: a bandwidth reservation protocol (called Helia [10]), a protocol to increase the path selection granularity to individual routers (called FABRID [11]), and a mechanism to guarantee connectivity to end hosts and services despite volumetric DDoS attacks (called Lightning Filter [3]). To ensure the secrecy of the transmitted data, the communication between sites is encrypted using VPN tunnels. In addition, we leverage a state-of-the-art DDoS defense solution called ACC-Turbo [12] to protect against pulse-wave DDoS attacks. Finally, we further improve privacy through traffic shaping and policing offered by a system called DITTO [13].

Figure 2 illustrates our solution, and in the following subsections, we provide more details about each of these components.

FIGURE 2: ILLUSTRATION OF VARIOUS THREATS AGAINST CONFIDENTIALITY, INTEGRITY, AND AVAILABILITY OF COMMUNICATIONS BETWEEN A SOURCE S AND A DESTINATION D IN A WAN DEPLOYMENT OVER PUBLIC INTERNET INFRASTRUCTURE, AND THE SOLUTIONS OUR ARCHITECTURE INCORPORATES TO MITIGATE THEM



B. SCION as the Foundation

SCION is a next-generation Internet architecture designed to improve security, scalability, and control over network traffic. SCION achieves path transparency by making the available (AS-level) forwarding paths visible to end hosts and allows path control by enabling end hosts to select among a set of offered paths. This improves reliability by allowing end hosts to quickly transition to alternative paths and can mitigate congestion through in-network multipath. SCION mitigates path hijacking attacks by design, since routing information is cryptographically secured in both the control and the data plane. Furthermore, control messages such as traceroute replies are authenticated to prevent off-path adversaries from modifying or injecting control messages. Based on the SCION architecture, researchers have developed several

mechanisms for more fine-grained path transparency and control as well as mechanisms to achieve minimum communication guarantees despite volumetric DDoS attacks. We leverage some of these extensions for our proposed WAN architecture. In the following sections, we provide more details about them.

C. Traffic Encryption with IPsec

IPsec [14] is a widely used protocol suite for encrypting and authenticating network traffic. IPsec can be used, among other things, to create encrypted and authenticated layer-3 tunnels between two endpoints. For these tunnels, the original IP packets are encrypted and encapsulated in a new IP packet. Therefore, only the IP addresses of the tunnel's endpoints are visible to potential eavesdroppers. We rely on IPsec tunnels to encrypt communications between WAN sites.

D. Intra-Domain Path Selection, Packet Source Authentication, and Path Validation with FABRID

FABRID extends SCION's capabilities of path transparency and control to the intra-domain router level. It is the first system that enables applications, and hence WAN operators, to forward traffic flexibly, potentially on multiple paths, selected to comply with user-defined preferences [11]. In FABRID, network operators communicate information about their internal router topologies in the form of router policies. These policies are then made accessible to applications along with the set of selectable forwarding paths. This allows each application to choose suitable router policies and encode the chosen policies within its data packets. Consequently, routers along these paths understand how to route traffic in accordance with the designated policies. These policies can encompass diverse router attributes within an AS, such as hardware specifications, geographic location, or manufacturer details. For instance, sensitive data may have to be sent exclusively through routers within specific jurisdictions. Likewise, services reliant on precise time synchronization might need the exclusive use of precision time protocol-capable routers. Some entities, such as governments or critical infrastructure operators, may mandate that traffic should avoid routers with known vulnerabilities or routers produced by untrusted manufacturers, resulting in paths comprised exclusively of recognized, trustworthy equipment. Apart from fine-tuning path transparency and control, FABRID enables on-path routers to cryptographically authenticate the source of each packet and extend the packets with proofs of transit, providing path validation for the source and destination hosts.

E. Bandwidth Reservation with Helia

Helia is a secure inter-domain bandwidth reservation protocol, allowing the dynamic allocation of bandwidth over SCION. It builds on the concept of flyover reservations, a fundamentally new approach for addressing the availability demands of critical low-volume applications. In contrast to path-based reservation systems, flyovers

are fine-grained “hop-based” bandwidth reservations on the level of individual ASs [10]. As Helia can offer forwarding guarantees despite large-scale volumetric DDoS attacks against network infrastructure, such that no off-path adversary can prevent the successful delivery of traffic sent over the reservation, Helia can be regarded as a cost-efficient alternative to the inherent guarantees of leased lines.

F. DDoS Protection with Lightning Filter and ACC-Turbo

DDoS attacks can target end hosts (e.g., web servers) or the network infrastructure (links or routers). For each of these targets, we propose a suitable defense system.

Lightning Filter [3] is a high-speed filtering system that protects (groups of) hosts or services from volumetric DDoS attacks. Just as Helia guarantees the successful forwarding of traffic through the network, Lightning Filter ensures access to the protected destination. Lightning Filter prevents DDoS attacks that rely on spoofed source IP addresses because it authenticates the source and payload of all incoming packets. Since currently deployed firewalls are often susceptible to such attacks, they can profit from Lightning Filter as a first layer of defense. Lightning Filter can thus reduce the packet drop rate under DDoS attacks by reducing the amount of traffic volume that has to be processed by specialized firewalls, and it can improve the accuracy of these firewalls by removing spoofed packets.

Volumetric DDoS attacks against routers and network links can be mitigated through bandwidth reservations, which provide fundamental availability guarantees. However, not all traffic is equally important, and therefore, a portion of traffic might be sent as best-effort traffic without bandwidth reservations. We therefore rely on ACC-Turbo [12] as an additional DDoS defense system. ACC-Turbo is highly effective in detecting and mitigating even pulse-wave DDoS attacks at line rate, making it the fastest in-network DDoS mitigation technique to date. Such in-network DDoS defense systems are significantly more effective in the absence of spoofed source addresses and thus benefit from FABRID’s source authentication.

G. Traffic Analysis Prevention with DITTO

Even though traffic is encrypted with IPsec tunnels between the WAN sites in our architecture, communication is still vulnerable to traffic analysis attacks. While FABRID’s fine-grained path control can mitigate such attacks to some degree, as it makes it possible to steer traffic around untrusted devices, this measure is insufficient if trust is not warranted. Furthermore, traffic analysis might occur not only at malicious or compromised devices but also at network links. Compared to on-path attackers actively delaying or dropping packets, detecting attackers performing traffic analysis is nearly impossible. We therefore rely on DITTO [13], a traffic obfuscation system to protect against traffic analysis. DITTO has been specifically designed to protect

links to WAN sites, scaling up to 100 Gbps links. In addition, it does not require modifications at the end hosts. DITTO adds padding and chaff packets at line rate, ensuring that outgoing traffic always follows a fixed pattern.

H. Summary

A combination of the abovementioned components results in a cost-efficient architecture to securely operate WANs over the public Internet. As summarized in Table II, this architecture implements comprehensive mitigations against all the threats discussed in Section 2. However, it is essential to note the potential drawback that traffic shaping and padding may increase the risk of congestion (Table I).

TABLE II: COMPONENTS OF OUR PROPOSED WAN ARCHITECTURE AND THE MITIGATIONS THEY IMPLEMENT

	Traffic encryption	Traffic shaping and padding	Traffic filtering	Traffic authentication	Path authentication	Path control	Traffic prioritization
IPsec	V			V			
DITTO		V					
Lightning Filter			V	V			V
FABRID				V		V	
Helia				V			V
SCION					V	V	
ACC-Turbo							V

4. THE ROAD TO DEPLOYMENT

Given that many of the components mentioned in Section 3 are still research prototypes, there is a gap between our proposed architecture and current ISP service offerings. To evaluate our architecture, we deployed a testbed across four WAN sites. In this section, we first describe our testbed and then focus on the technological readiness of each individual component. Afterwards, we discuss the remaining challenges and outline a possible use case.

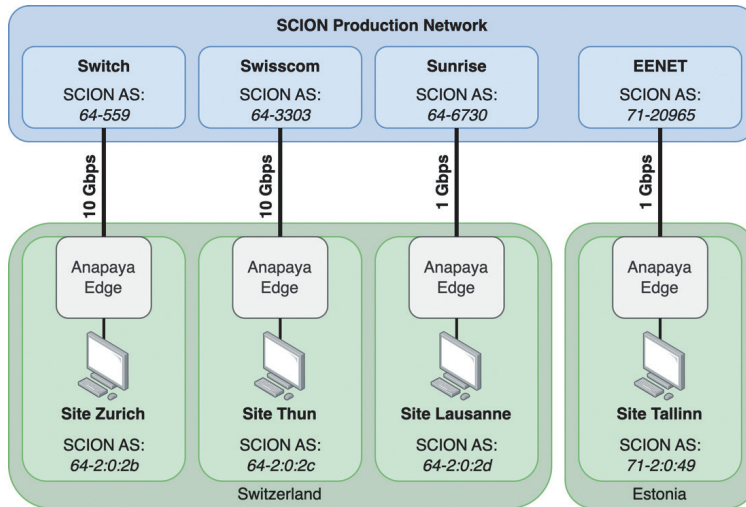
A. Testbed Deployment

We created a testbed that uses the SCION infrastructure provided by four different ISPs to establish a secure WAN connecting multiple locations. This deployment is primarily for experimental use, rather than incorporation into a production environment. Initially, we connected three separate locations across Switzerland to the SCION network over three different ISPs: Swisscom, Sunrise, and Switch. We expanded this deployment with an additional WAN node in Tallinn (Estonia) with EENet as the ISP. Each of these locations forms its own SCION AS.

Our setup relies on Anapaya Edge [15], a commercially available device running a SCION-IP gateway (SIG) and a SCION router. Anapaya is a spin-off of ETH Zurich, where large parts of SCION were developed. The SIG acts as a tunnel endpoint for IP packets, encapsulating them into SCION packets and delivering them to the specified SCION addresses. This way, end hosts and applications in our WAN deployment do not need to support SCION, as the SIG implements the whole protocol translation and path selection logic. Link access to the SCION network is provided through distinct ISPs that have established direct peering among each other.

In this SCION setup, all participants were assigned one or multiple AS identifiers, and the corresponding list is publicly available [16]. Figure 3 shows an overview of our current deployment.

FIGURE 3: OVERVIEW OF OUR SCION-BASED WAN DEPLOYMENT CONNECTING FOUR SITES IN TWO COUNTRIES



Our deployment allows us to direct all traffic between the sites natively over the SCION network. A major advantage of SCION is the availability of many paths, enabling the selection of various routes for routing traffic through a series of distinct hops. Figure 4 shows a subset of eight existing paths from the Anapaya Edge in Zurich to the Anapaya Edge in Thun. While SCION offers different AS-level paths, we can observe that there are also paths that traverse the same ASs and differ only with respect to their ingress and egress interfaces.

FIGURE 4: SUBSET OF AVAILABLE FORWARDING PATHS FROM ZURICH TO THUN AS SHOWN AT THE SIG IN ZURICH. PATHS ARE REPRESENTED AS AUTOMATIC SYSTEM SEQUENCES (HOPS) AND INTERFACE PAIRS. AN INTERFACE PAIR IS REPRESENTED AS “EG>IN,” WHERE “EG” IS THE EGRESS INTERFACE ID AND “IN” IS THE INGRESS INTERFACE ID

```

Available paths to 64-2:0:2c
4 Hops:
[0] Hops: [64-2:0:2b 1>24 64-559 17>1 64-3303 21>1 64-2:0:2c]
[1] Hops: [64-2:0:2b 1>24 64-559 17>1 64-3303 25>2 64-2:0:2c]
[2] Hops: [64-2:0:2b 1>24 64-559 19>9 64-3303 21>1 64-2:0:2c]
[3] Hops: [64-2:0:2b 1>24 64-559 19>9 64-3303 25>2 64-2:0:2c]
5 Hops:
[4] Hops: [64-2:0:2b 1>24 64-559 4>15 64-2:0:13 18>4 64-3303 21>1 64-2:0:2c]
[5] Hops: [64-2:0:2b 1>24 64-559 4>15 64-2:0:13 18>4 64-3303 25>2 64-2:0:2c]
[6] Hops: [64-2:0:2b 1>24 64-559 16>10 64-6730 11>11 64-3303 21>1 64-2:0:2c]
[7] Hops: [64-2:0:2b 1>24 64-559 16>10 64-6730 11>11 64-3303 25>2 64-2:0:2c]
[...]

```

We leveraged this level of path transparency to find all inter-domain paths offered by SCION at site Zurich. Table III shows the result of this evaluation: the total number of available paths, including paths that only differ in terms of their interfaces, from Zurich to the other sites in our deployment.

TABLE III: NUMBER OF AVAILABLE PATHS FROM ZURICH TO THUN, LAUSANNE, AND TALLINN. THE COLUMNS REPRESENT THE PATH LENGTH IN TERMS OF AS-LEVEL HOPS

Destination	Hops			
	4	5	6	7
Thun	4	32	66	8
Lausanne	2	14	35	-
Tallinn	2	-	-	-

We first verified the general connectivity between the sites using the SCION ping utility [17] without specifying a forwarding path, meaning that the SIG automatically

selected the path to each destination. The resulting round-trip time (RTT) and jitter are shown in Table IV. As expected, the RTT between sites in the same country was significantly lower than the RTT to the remote site in Tallinn.

TABLE IV: RTT AVERAGE AND STANDARD DEVIATION IN MILLISECONDS FOR 20 PROBE PACKETS FROM ZURICH AND TALLINN TO ALL OTHER WAN SITES. NOTE THAT THESE RESULTS DEPEND ON THE PATH SELECTED BY THE SIG, WHICH MAY CHANGE OVER TIME

Destination				
Source	Zurich	Thun	Lausanne	Tallinn
Zurich	-	12.08 (0.05)	4.89 (0.09)	59.42 (0.15)
Tallinn	60.09 (0.41)	72.77 (0.22)	65.12 (0.33)	-

To evaluate the impact of different inter-domain paths on the RTT, we measured the RTT between Zurich and Thun using various paths of different lengths. This time, we used the SCION ping utility to explicitly specify the forwarding path. We manually selected a total of 12 paths, three each of lengths four, five, six, and seven. Table V shows the obtained RTT and jitter values. The measurements show that paths of the same length can vary significantly in terms of their RTT. At the same time, longer paths do not necessarily imply a higher RTT, as can be observed for the values up to six AS-level hops, where the RTT values are similar. We observe consistently low jitter irrespective of the forwarding path.

TABLE V: RTT AVERAGE AND STANDARD DEVIATION IN MILLISECONDS FROM ZURICH TO THUN FOR PATHS WITH DIFFERENT NUMBERS OF AS-LEVEL HOPS, AVERAGING OVER 20 PROBES. EACH COLUMN REFERS TO A DIFFERENT PATH

	4 hops			5 hops			6 hops			7 hops		
RTT (ms)	5.4	9.3	11.7	5.7	9.5	12.0	6.2	9.6	15.3	11.7	18.6	22.1
RTT standard deviation (ms)	0.04	0.08	0.06	0.44	0.03	0.07	0.06	0.52	0.09	0.07	1.10	0.19

B. Readiness of the Proposed Components

Technology readiness level (TRL) is a widely used metric to describe the maturity of a technology [18]. It uses a scale from TRL 1, which applies to technology whose basic principles have been observed and reported, to TRL 9, which applies to technology that is used in an actual system in production.

For each component of our proposed architecture, Table VI indicates the current TRL and whether there exist ISPs that readily offer it as a service. Most technologies are in a rather early stage (TRL 3)—this is because they were only recently proposed by academic research. For some components, ISP support is not needed because they can be implemented end-to-end by the WAN operator without any ISP integration. SCION is fully supported by the four commercial ISPs in our testbed and offered as an experimental service in their operational environment (TRL 7).

TABLE VI: TRLS OF THE COMPONENTS COMPRISING OUR PROPOSED WAN ARCHITECTURE

Technology	Offered by ISPs	TRL
IPsec	Not needed	9 (actual system proven in operational environment)
SCION connectivity	Yes	7 (system prototype demonstration in operational environment)
FABRID	Not yet	3 (experimental proof of concept)
Helia	Not yet	3
Lightning Filter	Not yet	3
ACC-Turbo	Not yet	3
DITTO	Not needed	3

C. Remaining Challenges

There are several challenges in implementing and deploying our proposed WAN architecture. The biggest of these is the lack of support from commercial ISPs for the needed security components (see Table VI). The exception is SCION connectivity, which is already provided by commercial ISPs. However, SCION is not yet available in all countries, meaning that it may not be possible to connect some WAN sites to the production network over an arbitrary ISP. Still, the current SCION deployment readily covers networks in Europe, Asia, and North America and is rapidly expanding [16]. There are also challenges regarding differing hardware requirements; the different security solutions have only been evaluated independently of each other so far. ACC-Turbo and DITTO have been implemented in P4-compatible devices. Today, these devices do not natively support cryptographic algorithms such as Advanced Encryption Standard (AES). However, AES is required by SCION and its extensions. Recently, a P4 router implementation has been presented that uses an accelerator for the cryptographic validations and can thus forward SCION traffic on the Intel Tofino 2 [19] at a rate of more than 3 terabits per second [20]. Given these advancements, a P4 implementation of Helia, FABRID, and Lightning Filter, each requiring further

AES key expansions and block cipher computations compared to SCION, might soon become feasible. Nevertheless, the components of our proposed WAN solution can also be deployed on other platforms; for example, some components have been implemented and evaluated using Intel’s Data Plane Development Kit (DPDK). In financial terms, precisely estimating or predicting the costs of our proposed architecture is challenging. Nevertheless, we expect the costs to be significantly lower compared to solutions such as leased lines, as our architecture operates on an existing, cost-efficient public Internet fabric and existing SCION routers used in production rely on DPDK, therefore new security mechanisms can be installed through software updates. Lastly, while every component has already been independently analyzed for its security properties, the security of the entire system has not been explored—an interesting and important avenue for future work.

D. Use Case: Multinational Collaboration


Our architecture is suitable for many types of WANs, from small companies to large multinational organizations with sites across the globe. Besides our own deployment, another interesting use case could be secure communications among different countries, for example, transmitting sensitive data between their data centers, where this traffic should ideally only traverse other member countries.

As such, this scenario has similarities with the collaboration among NATO countries, which might want to securely interconnect their data centers or cyber ranges for training purposes [21].

As illustrated in Figure 5, our proposed architecture constitutes an ideal solution for the following reasons:

- The individual sites can be operated independently by the respective countries.
- WAN connections use existing infrastructure; therefore, no additional cables need to be installed.
- Path control makes it possible to ensure that traffic only crosses member countries, even if it is not the shortest path.
- If there is no path that traverses only member countries, the number of non-member countries can be minimized. Encryption and traffic shaping minimize the attack surface in this case.

FIGURE 5: WAN CONNECTING DATA CENTERS LOCATED IN DIFFERENT COUNTRIES. OUR ARCHITECTURE ALLOWS ROUTING TRAFFIC ONLY ALONG PATHS THAT DO NOT LEAVE MEMBER COUNTRIES IF THIS IS POSSIBLE (A, B). IF IT IS NOT POSSIBLE (C), THE NUMBER OF TRAVERSED NON-MEMBER COUNTRIES CAN BE MINIMIZED AND OTHER MITIGATION MEASURES LIMIT THE RISK OF A SUCCESSFUL ATTACK



5. CONCLUSION

In this paper, we analyzed the challenges in designing secure WANs for large organizations to interconnect their geographically distributed sites. We found that such a WAN can be built on the shared infrastructure of public ISPs thanks to the increasing adoption of the SCION next-generation Internet architecture and recently introduced technologies. Our proposed solution leverages SCION and recent results from the research community to achieve strong security guarantees while significantly reducing costs compared to the currently used WAN approaches that are based on leased lines.

To verify the feasibility of this approach, we implemented and evaluated basic SCION connectivity at multiple WAN sites in two countries. In future work, we plan to extend this testbed both geographically by adding more sites and technologically by implementing and deploying the missing components.

Once comprehensively implemented and deployed, the proposed architecture would allow organizations of any size to build secure WANs over the public Internet.

REFERENCES

- [1] Z. Yang, Y. Cui, B. Li, Y. Liu, and Y. Xu, "Software-defined wide area network (SD-WAN): Architecture, advances and opportunities," in *Proceedings of the International Conference on Computer Communications and Networks (ICCCN)*, Jul. 2019, doi: 10.1109/icccn.2019.8847124.

- [2] P. Zilberman, R. Puzis, and Y. Elovici, "On network footprint of traffic inspection and filtering at global scrubbing centers," *IEEE Trans. Dependable Secure Comput.*, vol. 14, no. 5, 2015.
- [3] L. Chuat et al., *The Complete Guide to SCION*. Springer International Publishing, 2022, doi: 10.1007/978-3-031-05288-0.
- [4] E. Papadogiannaki and S. Ioannidis, "A survey on encrypted network traffic analysis applications, techniques, and countermeasures," *ACM Comput. Surv.*, vol. 54, no. 6, Jul. 2021, doi: 10.1145/3457904.
- [5] G. Schmid, "Thirty years of DNS insecurity: Current issues and perspectives," *IEEE Commun. Surv. Tutor.*, vol. 23, no. 4, 2021, doi: 10.1109/comst.2021.3105741.
- [6] "DDoS threat intelligence report: Issue 11." NETSCOUT. Dec. 2023. [Online]. Available: <https://perma.cc/V44P-543K>
- [7] O. Yoachimik. "Cloudflare DDoS threat report for 2022 Q4." Cloudflare Blog. Oct. 2023. [Online]. Available: <https://perma.cc/Q6VB-BG2J>
- [8] "DDoS threat report for 2023 Q3." Cloudflare Blog. Accessed: Jan. 4, 2024. [Online]. Available: <https://blog.cloudflare.com/ddos-threat-report-2023-q3>
- [9] C. Xu, J. Zhao, and G.-M. Muntean, "Congestion control design for multipath transport protocols: A survey," *IEEE Commun. Surv. Tutor.*, vol. 18, no. 4, 2016, doi: 10.1109/comst.2016.2558818.
- [10] M. Wyss, G. Giuliani, J. Mohler, and A. Perrig, "Protecting critical inter-domain communication through flyover reservations," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, Nov. 2022, doi: 10.1145/3548606.3560582.
- [11] C. Krähenbühl, M. Wyss, D. Basin, V. Lenders, A. Perrig, and M. Strohmeier, "FABRID: Flexible attestation-based routing for inter-domain networks," in *Proceedings of the USENIX Security Symposium*, 2023.
- [12] A. G. Alcoz, M. Strohmeier, V. Lenders, and L. Vanbever, "Aggregate-based congestion control for pulse-wave DDoS defense," in *Proceedings of the ACM SIGCOMM Conference*, 2022, doi: 10.1145/3544216.3544263.
- [13] R. Meier, V. Lenders, and L. Vanbever, "Ditto: WAN traffic obfuscation at line rate," in *Proceedings of the Symposium on Network and Distributed Systems Security (NDSS)*, 2022, doi: 10.14722/ndss.2022.24056.
- [14] "IP security protocol (IPsec)." IETF Datatracker. Accessed: Dec. 29, 2023. [Online]. Available: <https://datatracker.ietf.org/wg/ipsec/>
- [15] "Anapaya Edge overview." Anapaya. Accessed: Apr. 4, 2024. [Online]. Available: <https://docs.anapaya.net/en/latest/edge/overview/>
- [16] "ISD and AS assignments." Anapaya. Accessed: Jan. 4, 2024. [Online]. Available: <https://docs.anapaya.net/en/latest/resources/isd-as-assignments/>
- [17] "SCION ping." Anapaya. Accessed: Jan. 5, 2024. [Online]. Available: https://scion.docs.anapaya.net/en/latest/command/scion/scion_ping.html
- [18] "Technology readiness levels." NASA. Accessed: Jan. 5, 2024. [Online]. Available: <https://www.nasa.gov/directorates/somd/space-communications-navigation-program/technology-readiness-levels/>
- [19] "Intel Tofino 2." Intel. Accessed: Jan. 4, 2024. [Online]. Available: <https://www.intel.com/content/www/us/en/products/details/network-io/intelligent-fabric-processors/tofino-2.html>
- [20] L.-C. Schulz, R. Wehner, and D. Hausheer, "Cryptographic path validation for SCION in P4," in *Proceedings of the 6th on European P4 Workshop*, 2023.
- [21] "Cyber Ranges Federation Project reaches new milestone." European Defence Agency. Accessed: Jan. 4, 2024. [Online]. Available: <https://eda.europa.eu/news-and-events/news/2018/09/13/cyber-ranges-federation-project-reaches-new-milestone>