Routing Gets Personal: Welcome to /32 to the client

Bachelor Report

Department of Computer Science
OST – University of Applied Sciences
Campus Rapperswil-Jona

Spring Term 2025

Author	Stefan Meyer
Advisors	Urs Baumann und Jan Untersander
External Co-Examiner	Philip Schmid
Internal Co-Examiner	Prof. Stefan F. Keller

Contents

1. Abstract

No matter which network architecture or topology is used, large Layer 2 domains can be found everywhere. They are also used in campus networks to connect end devices and segment them appropriately into their own broadcast domains. However, using extended L2 domains also brings all the restrictions, disadvantages and behaviour patterns belonging to this layer into the network environment. These include the limitations of the Spanning Tree Protocol, the poor scalability of L2 networks, and the flooding of broadcast, unknown unicast and multicast (BUM) network traffic. The bachelor's thesis investigates a novel approach to circumvent these issues, referred to as the «/32 environment», which aims to minimize the size of Layer 2 domains as much as possible. This is achieved by assigning a /32 subnet mask to each end device alongside its IP address. Consequently, each end device operates within its own network, resulting in a Layer 2 domain containing only a single IP address. This solution follows the idea of extending Layer 3 down to the end device. The approach was validated by establishing a physical test environment (Cisco) and a virtual test environment based on «containerlab» (Arista). It gave the possibility to test various scenarios and network protocols. These experiments have verified the technical feasibility of this approach and its associated advantages. In addition, all tested protocols perform correctly in a /32 environment. Currently, however, this approach cannot meet the requirements of a campus network as the router OS from Arista and Cisco either has bugs or lacks the necessary functions. Future studies could examine how such an implementation would perform in WLAN infrastructures or data centres, and the challenges that might arise.

2. Management Summary

In the following chapters the content of the management summary can be found.

2.1 Introduction

Regardless of the network architecture or topology in use: Layer 2 networks can be found everywhere. Layer 2 is particularly widespread in the access layer, where L2 switches are used for the most part, as all kinds of end devices are connected to the network, whether by cable or wirelessly. However, the use of large Layer 2 domains also brings disadvantages and limitations to the network, which can lead to minor or major problems depending on the situation.

In order to shrink these Layer 2 domains as much as possible in the network (to point-to-point), a new approach is being taken where each end device is on its own network with a /32 subnet mask. This means that each device is treated as a separate network and is routed to Layer 3 accordingly. This modification could make the entire network infrastructure less prone to faults and more efficient because everything is based on routing. This new approach is referred to in this bachelor’s thesis by the term «/32 environment».

2.2 Approach and technology

In a first step, the disadvantages of Layer 2 domains were discussed in more detail to emphasize the motivation behind this new approach for a /32 environment. The focus of this work and the tests are based on a campus network. For this reason, it was analyzed which protocols and services

mainly occur in such an environment. Based on these findings, a selection of protocols was chosen that will be tested in detail for their functionality in the /32 test environments. To round off the overview, attention was also drawn to protocols and device groups that would no longer function correctly in an exclusive Layer 3 to the end device approach.

With these initial insights, two identical test environments with two different network device manufacturers were set up in the practical part of this work. Arista was used in a virtual test lab, which is based on «containerlab» and works exclusively with Docker containers (Figure 1). Cisco was used in the physical test lab to be able to cover the real-world part. The /32 network was then set up and configured accordingly based on these test lab infrastructures. The problems encountered during this process, as well as solutions and workarounds that were discovered, were recorded in detail and investigated further where necessary. The same approach as for IPv4 with /32 was then followed and tested in an IPv6 environment with a /128.

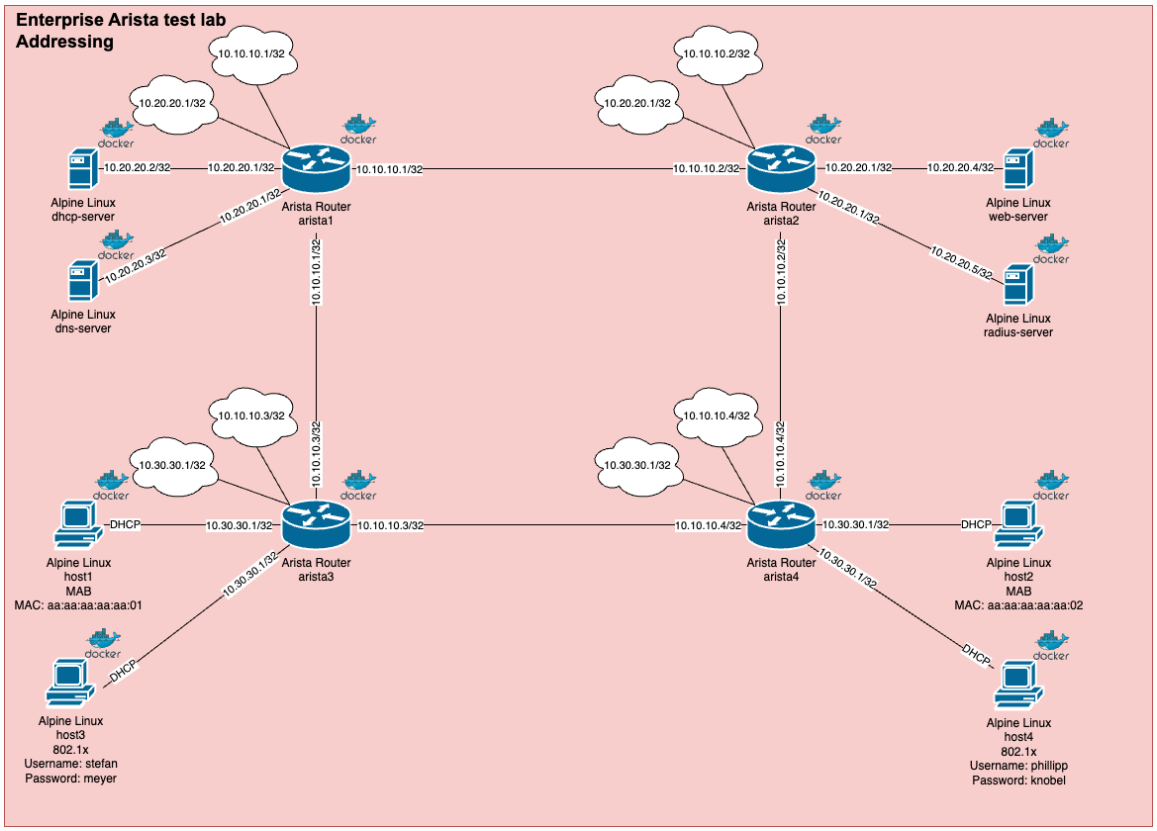

Figure 1: Network diagram enterprise Arista test lab

2.3 Results and outlook

This work successfully demonstrated that a /32 approach can be implemented successfully and with the expected benefits from a purely technical perspective. Based on the two support matrices (one of those in Figure 2) that were created during the work, two further results came to light. The first one is that, as of today, it is not possible to set up a correctly and dynamically functioning /32 environment network with Arista or Cisco, as both manufacturers have restrictions or bugs for which no solutions or workarounds are available. The second result confirms that all protocols work correctly in a /32 environment.

Some protocols and devices will still require large Layer 2 networks. However, this work demonstrates that where they aren't necessary, Layer 3 can be expanded up to the end device as long as router vendors provide the necessary features and bug fixes.

Topic	Arista	Cisco
Set IP address with /32 subnet mask on routed interface		
Routing works with Anycast Gateway		
DHCP Option 82 works as intended		
Automatically set static host route for DHCP host		
More than one DHCP host can be connected to the same router with the same Anycast Gateway		
Routing works with DHCP bindings		
Set interface profile / template with Radius Access-Accept message		
Interface profile / template does support IP related commands		
MAB can be set on routed interface without error message		
MAB configuration works on routed interface		
802.1x configuration behaves normally on routed interface		

Figure 2: Support matrix /32 environment Arista and Cisco - Green = Works, Orange = Works partially, Red = Does not work

3. Introduction

This chapter mainly deals with organizational aspects.

3.1 Assignment

The goal of the work is to use appropriate test labs to find out whether a /32 environment is really feasible or practicable, and whether it can be implemented independently of the manufacturer. It also needs to be clarified if and how such a network can be operated and what the implications are for the most used protocols. In addition, device groups and (enterprise) functions should be identified that would not work at all or only partially with this new approach. Moreover, attention will be drawn to limitations. Further explanations will be given as to why the /32 approach does not work or should be pursued in the future.

This bachelor's thesis is a feasibility study. The following is a rough summary of the deliverables that are expected at the end:

- A small-scale prototype or simulation of a /32-based network
- Evaluation of the behavior, performance, and manageability of the approach
- A detailed written report according to the requirements
- Well-structured documentation of the prototype or testbed setup, including instructions for reproducing key results
- A summary of lessons learned, open questions, and potential directions for future work in this area

3.2 Structure bachelor's thesis

For reasons of clarity and to be able to forward the findings to the manufacturers Arista and Cisco, two separate documents have been created. The advisors have agreed upon this.

Brief overview of the two documents:

- **Bachelor Report**: On one hand, the administrative topics of the bachelor's thesis are described, and on the other hand, the topic of /32 network environments are introduced. It also provides an overview of why this work is relevant and what the current problems are today's networks face. Various protocol groups in a campus network are also discussed. Based on these findings possible problems in /32 environments are pointed out. In addition, a table of important protocols has been defined that must be tested in the /32 network environment. Finally, the results and limitations of the extensive tests in the test labs of Arista and Cisco are summarized and a conclusion is formulated.
- **Lab Documentation**: This document is mainly about the practical part of the bachelor's thesis. At the beginning the various test environments used for the numerous tests are described. Furthermore, the extensive and detailed tests show what works in the /32 network environments with the hardware manufacturers Arista and Cisco and what works only partially or not at all. Moreover, the various protocols and services that were defined in advance in the «Bachelor Report» are tested here. The document is completed with the diagrams, topologies and the configuration files used. This document has been designed so that it can be sent to Arista, Cisco or other external persons to inform them about the problems and bugs found.

3.3 Formalities

This bachelor's thesis with the title «Routing Gets Personal: Welcome to /32 to the client» falls under the subject area of «Network and Cloud Infrastructure». The time budget for a bachelor's thesis is 360 hours and includes 16 semester weeks and a block week in which you can work fully on the thesis. The ECTS credit for this type of work is 12 ECTS.

3.4 Risk analysis

Since this is a feasibility study, no risk analysis will be prepared and documented. This has been agreed upon by the advisors.

3.5 Scope

As this bachelor's thesis is limited in time and covers a large number of different network topics that need to be processed, the scope must be clearly defined. Research is used to define which protocols and services are to be tested in a /32 environment. They are limited to the most popular and important ones that occur in a campus network. The device categories used are limited to tests based on a client-server relationship, covering the macOS, Windows and Linux operating systems. As there is simply not enough time available to build up good expertise or the technical equipment and resources are not available, the whole area of OT devices and OT networks is only discussed theoretically and not verified with practical tests. The entire test environment will exclusively focus on wired connections (WLAN is not covered) and include only the most necessary configurations.

3.6 Motivation

As mentioned in the previous chapter, all of today's networks are in some form based on Layer 2 domains, either because it has evolved from previous best practice recommendations or because the correct operation of certain device groups and protocols depends on them.

In the following chapter, the disadvantages and challenges of large Layer 2 domains are discussed in order to underline why this bachelor's thesis is being written and why this topic is being examined in more detail.

3.6.1 Layer 2 challenges and problems

With the use of Layer 2 domains, you also have to deal with their restrictions, disadvantages and behaviour patterns. The following list explains the reasons why network engineers do not like Layer 2:

- Layer 2 networks require the use of the Spanning Tree Protocol (STP) to establish a loop-free topology. The use of STP results in the following limitations [1]:
    - It is not possible to utilise the full performance and possibilities of the network because STP blocks redundant paths to maintain a loop-free network. However, this behaviour blocks available links, which leads to a loss of bandwidth and capacity.
    - Direct and optimal paths in the network could be blocked if the root bridge is not selected correctly or if some misconfiguration or default settings are applied to the network devices. This can lead to suboptimal paths, resulting in higher latency and reduced network performance.
    - In case of an interface or switch failure the whole topology needs to be recalutated with the goal of a loop-free network. This recalculation process takes time and causes a high convergence time. This delay has a negative impact on time-sensitive applications like VoIP and high-bandwidth interfaces.
    - Equal-Cost Multi-Path (ECMP) is not supported with STP which results in a loss of available bandwidth. Again, the reason is because redundant links are blocked to create a loop-free topology.
    - The goal of STP is to prevent loops in a Layer 2 topology. But in certain scenarios or under certain precondition, loops can still occur. When traffic loops endlessly, we refer to it as a broadcast storm. This can happen despite correctly configured STP. Triggers can be software bugs, faulty network interfaces or targeted malicious attacks. The traffic will loop indefinitely until the network goes down, because the hardware is unable to handle this load anymore. This is also referred to as a network meltdown. Forwarded loops can also affect core interfaces, potentially causing the entire network to crash due to the absence of a comparable solution like TTL at Layer 3 [2].
    - Dual-Homing is not supported with STP without some additional technology. Servers or hosts connected to two switch interfaces simultaneously can only utilize one link. If this link fails, the traffic flow is disrupted until the STP recalculation process is completed.
    - The challenges that lie in the design of STP (various STP variants exist) and additional, not fully finalised implementations such as protection mechanisms (e.g.

BPDUguard [3] and storm control [4]) within STP also lead repeatedly to problems and, in the worst case, to a network meltdown.

- Every L2 domain (VLAN) is a single failure domain primarily due to BUM (broadcast, unknown unicast, and multicast frames) flooding [4].
- L2 domains do not scale well, because the larger the network becomes, the more traffic is sent and the higher the risk that the L2 network will be overwhelmed. BUM traffic is particularly important here, because this type of traffic is distributed throughout the entire L2 network. Every broadcast that is sent in the L2 network must be processed by every host and network device. In addition, the BUM traffic also requires bandwidth on each link [5]. If you want to create L2 networks that are as small as possible and less vulnerable for these reasons, you will come up against the limitation that you can only create a maximum of 4094 different VLANs [6]. This is not enough, especially in data centers.
- Software errors or malwares and viruses that result in uncontrolled flooding of a host or server affect all other hosts in the same L2 domain. In addition, such flooding impacts the CPU load of L2 switches and routers that have the corresponding IP address configured in this L2 domain [5].

Of course, there are solutions for some of the L2 problems listed above. However, these have further limitations and complicate the entire network configuration and troubleshooting.

3.6.2 Solution approaches

Due to the reduced reliability in Layer 2, companies with highly available applications had to distribute their applications across several different Layer 2 domains to mitigate the single failure domain.

Later, standards such as TRILL and SPB should at least replace STP with its weaknesses in Layer 2. Both are based on the IS-IS routing protocol for determining optimal connection paths (L2 routing). With both standards, links are no longer blocked and all available connections are used. Multipath routing is also supported [7]. However, both standards are not widely used and if they are, they are almost exclusively found in data centers. A main reason for this development was that many manufacturers created their own proprietary protocols based on TRILL (Cisco with FabricPath, Brocade with VCS), which were not compatible with each other and required the corresponding hardware. In addition, there were many incompatibilities that occurred between different manufacturers regardless of the standard and the industry rejected the premise [8].

With the widespread use of VXLAN, which also became established as a standard, Layer 2 packets can be tunneled via Layer 3 networks. VXLAN is supported by many manufacturers and has been accepted by the industry. But here too, only the STP issue has been resolved. Extended Layer 2 networks are still used, with all their weaknesses and characteristics. For example, the BUM traffic is still present.

4. Research

This chapter is about collecting information on similar approaches to a /32 environment. It also analyzes which protocols occur in a campus network and highlights problems associated with /32 environments.

4.1 Related work and findings

Nothing could be found regarding the approach of using a /32 environment exclusively and the idea of extending Layer 3 down to the end device.

Solutions such as TRILL, SPB and VXLAN [9] have been developed (3.6.2) to solve the problems and limitations of STP described above (3.6.1). However, these solutions still explicitly use extended Layer 2 networks, which does not eliminate the disadvantages of this Layer.

4.2 Traffic in a campus network

To find out whether a /32 network environment can also be utilized in a campus network [10], you need to know what kind of traffic is primarily found in such a network.

The network world has seen the rise of two different types of network traffic [11]:

- East-West traffic
- North-South traffic

East-West traffic refers to network traffic that occurs within the same network scope. This type of communication takes place and stays within the network itself. Examples of East-West traffic include communication between servers in a datacenter belonging to the same network scope or the distribution of STP information between switches.
On the other hand, North-South traffic refers to traffic that flows in and out of a network scope. For instance, client-to-server traffic is an example of this type of traffic. In this scenario, traffic from the clients leaves the client network scope and enters the server network scope where the servers are located.
North-South traffic, per Cisco's Global Cloud Index from 2014, dominates campus networks, comprising over 90% of the traffic [12]. This traffic includes activities such as internet access, cloud service utilization and access to servers. This value is unlikely to change any time soon, as the most data is still exchanged via a classic client-to-server connection [11]. However, with the increasing use of IoT and OT devices, the proportion of East-West traffic could increase, as these devices communicate directly with each other within the same network scope.

In data centers, the East-West traffic is more distributed, comprising approximately 76% of the total traffic [12]. This trend has been further reinforced by the widespread use of virtualization. But even in the data center sector, North-South traffic will sooner or later overtake East-West traffic. This as a consequence of the trend towards public clouds. As a result, traffic is no longer limited to the data center network scope but flows freely in and out of the data center to access the public cloud [13].

4.2.1 Traffic categories in a campus network

The studies found that deal with the data traffic of a campus network only analyze the data traffic between the campus network and the Internet. Unfortunately, no documented information could be

found regarding data traffic and the most common protocols within a campus network. This is likely due to data protection concerns, as no one wants to publish sensitive data.

The findings and results of the above-mentioned studies are briefly summarized below. As mentioned, only the traffic leaving or entering the campus network was analyzed.

- A study from 2011 [14] analyzed the bandwidth consumption of a mid-size university for an entire year to find out which protocols or applications require the most bandwidth. The following are the most common protocols which were listed with reference to a relevant bandwidth utilization: HTTP, HTTPS, SSH, RTMP, IPSEC-ESP und SMB.
- In a research conducted in 2020 [15], a K-means clustering algorithm was used to analyze user internet access patterns and identify network trends. The study was based on data which was collected for two days from a campus network. At the end the protocols DNS, HTTP, HTTPS, MySQL, SSH, NTP and Telnet occurred the most.
- In a 2017 study [16], the data from the edge routers of the campus network from the University of Calgary were recorded over a period of seven weeks. The goal was to characterize and identify period traffic. The overall traffic on a per-connection basis is based on 73% TCP, 23% UDP and 4% ICMP. As was to be expected, HTTP, HTTPS and DNS are the most popular protocols. A summary of the top 10 ports and protocols can be found below (Figure 3).

Number of Connections	Protocol	Port	Registered Service
3,001,962,604	TCP	443	HTTPS
2,949,772,466	TCP	23	Telnet
2,144,724,698	UDP	53	DNS
1,892,680,004	TCP	80	HTTP
485,566,113	ICMP	0	ICMP Network Unreachable
376,924,831	TCP	22	SSH
328,420,663	TCP	5358	WSD
272,396,683	TCP	7547	CPE WAN
222,413,478	TCP	2323	Telnet (alt port)

Figure 3: Summary of the top 10 ports and protocols

The popularity of these non-standard protocols (e.g. Telnet, WSD and CPE WAN) is likely due to malware seeking out vulnerable machines to exploit.

A collection of protocols that occur in the internal network traffic of a campus network has been compiled from several sources on the Internet [17], [18], [19] and from own experience.

The protocols can be divided into different categories, which are summarized in the following tables.

4.2.1.1 Client traffic (user-generated traffic)

This category includes traffic generated by users using network services (Table 1).

Traffic type	Protocols used	Examples
Web browsing	HTTP, HTTPS	Websites
Email communication	SMTP, IMAP, POP3	Receive and sending emails
File transfers	FTP, SFTP, NFS, SMB, WebDAV	Upload and download of files
Streaming Media	RTP, RTSP	Live video streaming
Communication tools	SIP, WebRTC	Voice over IP
Remote Access	SSH, RDP, IPSec, IKEv2, OpenVPN	VPN solutions and remote access to servers

Table 1: Client traffic (user-generated traffic)

4.2.1.2 Server traffic

This category includes the types of traffic that are most common between servers (Table 2).

Traffic type	Protocols used	Examples
Database queries	MySQL, PostgreSQL	Database access
Authentication and directory services	LDAP, Kerberos, RADIUS, TACACS+	User authentication
File sharing and storage	NFS, CIFS, iSCSI	Shared drives and backups
Printing services	IPP, SMB, AirPrint	Printing systems
Email servers	SMTP, IMAP, POP3	Receive and sending emails
Application servers	HTTP, HTTPS	Web applications

Table 2: Server traffic

4.2.1.3 Network infrastructure traffic

This category includes the types of traffic that are shared inside a network infrastructure (Table 3).

Traffic type	Protocols used	Examples
Network management	ARP, DNS, DHCP, TCP, UDP, IGMP, Apple Bonjour	Establishing connection between devices
Switching traffic	VLAN (802.1Q), STP, RSTP, MSTP, VTP	Layer 2 networking
Routing protocols	OSPF, BGP, EIGRP, RIP	Layer 3 networking
Network monitoring	SNMP, ICMP, Syslog	Device monitoring and logs
Time synchronization	NTP, NTPS	Synchronize device clocks

Table 3: Network infrastructure traffic

4.2.1.4 Security and access control traffic

This category includes the types of traffic which are in charge to enforce and monitor security aspects of a network (Table 4).

Traffic type	Protocols used	Examples
Access control	802.1x	Device and user authentication
Monitoring and logging	Syslog	Logging of events

Table 4: Security and access control traffic

4.2.1.5 IoT and smart device traffic

The final category is dedicated to the group of IoT devices that are being used increasingly (Table 5).

Traffic type	Protocols used	Examples
Building automation	KNX, Matter, BACnet/IP	Lighting and ventilation control
Security systems	RTSP, MQTT	Videocamera and door controls
Environmental sensors	CoAP, LWM2M	Temperature and humidity sensors

Table 5: IoT and smart device traffic

4.3 Testing protocols in /32 environment

As it is not possible to test all the protocols and services mentioned above within the given time, we test a selection of the most important and most frequently used protocols. This means that these protocols are checked for correct functioning within the Testing Lab.

In a /32 environment, each end device operates within its own network, which has a single IP address that can lead to challenging phenomena and problems.

Table 6 lists the protocols that are specifically analyzed and tested. A comment is used to indicate whether or not the protocol will work in a /32 environment with the current level of knowledge. This assumption is then verified later with appropriate tests.

Protocol	Hypothesis
DHCP	DHCP discover uses the broadcast IP address 255.255.255.255 to search for a DHCP server. With the configuration of a DHCP relay the Layer 2 boundaries can be crossed.
DNS	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.
HTTP, HTTPS	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.
MAB	It depends on whether MAB can be configured on a non-switching interface. In the background of MAB is RADIUS which operates on the Application Layer.

802.1x	It depends on whether 802.1x can be configured on a non-switching interface. In the background of 802.1x is RADIUS which operates on the Application Layer.
ARP [20]	Is still needed in Layer 2 point-to-point networks. For example, the MAC address of the default gateway is searched from the host via ARP. ARP uses the broadcast MAC address of FF:FF:FF:FF:FF:FF and is limited to the same Layer 2 domain.
SSH	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.
RDP	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.
SMB	Should operate without problems, because works on Session Layer and requires only Layer 3 routing.
SNMP	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.
NTP	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.
SIP	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.
RTP	Should operate without problems, because works on Application Layer and requires only Layer 3 routing.

Table 6: Protocols which are tested in a /32 environment

4.4 Transmitting methods

Whether a protocol or service functions correctly in a /32 environment also depends on the transmitting method on which it is based. All L3 variants of the transmitting methods are unaffected by this new /32 approach and will continue to work. Consequently, only the L2 variants of these transmitting methods are discussed below.

4.4.1 Unicast Layer 2

Unicasts on Layer 2 are used when a direct peer-to-peer connection between two devices in the same Layer 2 domain is required. The majority of protocols and services are able to work on Layer 2 and Layer 3. Nevertheless, there are a few protocols that are based exclusively on Layer 2 unicasts. Two of these are shown in the provided table (Table 7), but it's important to note that these are classic datacenter protocols, not campus protocols.

Application	Protocol
Network protocol that allows block storage devices to be accessed over an Ethernet network [21]	ATA over Ethernet (AoE)
Network protocol that allows Fibre Channel (FC) data traffic to be transmitted over standard Ethernet networks [22]	Fibre Channel over Ethernet (FCoE)

Table 7: Protocols which rely on L2 unicast

4.4.2 Multicast Layer 2

Layer 2 multicasts are primarily used for the efficient distribution of messages to which the corresponding protocols and services have subscribed. However, multicast is also frequently used to find services and devices in the same network. Another important point is the checking of device health information, for example in connection with services that maintain the redundancy of important devices.

Table 8 is a non-exhaustive list of examples that are dependent on L2 multicasts. It should be mentioned that many services and protocols also have the option of connecting via Layer 3. It is also possible to set up «multicast policies [23]» or IGMP proxies / mDNS reflectors that forward Layer 2 multicasts to other networks. However, this is not according to the RFC standard and can therefore lead to problems.

Application	Protocol
Network message used by the STP to exchange information between switches	BPDU
Network protocol that increases the availability of the default gateway	VRRP
Cisco protocol that provides network redundancy with multiple routers	HSRP
IPTV (Internet Protocol Television)	IGMP
Apple AirPlay, Apple AirPrint [24]	mDNS, Bonjour
Google Chromecast	mDNS, SSDP
Smart Home Assistants (Amazon Alexa, Google Assistant, Apple HomeKit)	mDNS, SSDP
Smart Lighting & Automation (Philips Hue, Zigbee Bridges)	mDNS, UPnP

Table 8: Protocols which rely on L2 multicast

4.4.3 Broadcast Layer 2

The aim of broadcasts is to find devices / services and the distribution of information. In contrast to a multicast, with a broadcast the message is sent to all devices in the same L2 network. However, only the devices for which this message is intended process it. Well-known protocols are ARP and DHCP, which are based on this type of broadcasts.

Layer 2 broadcasts are also necessary to make «silent hosts [25]» «visible» again in the network. This category of hosts does not send regular traffic and therefore the MAC address is deleted from the MAC address table at the switch interface. As a result, these devices may no longer be accessible. Most of these devices are very old, but various IoT devices, industrial machines and building automation systems such as door locks, sensors or heaters are also known for such behavior.

The following is a non-exhaustive list of examples that work with L2 broadcasts (Table 9).

Application	Protocol
IP Address Resolution	ARP
IPv4 Address Auto-Configuration	DHCP
Network Device Discovery (HP, Cisco, etc.)	CDP, LLDP
Wake-on-LAN	Magic Packet
Cisco WLC Discovery [26]	LWAPP Discovery Request
IoT & Smart Home Devices	mDNS
Industrial Automation (PLC, BACnet/IP, SCADA)	BACnet/IP, PROFINET
Building Automation (HVAC, Access Control, etc.)	BACnet/IP

Table 9: Protocols which rely on L2 broadcast

4.5 Affected services and features

As mentioned above (4.4), all services and devices that rely on Layer 2 will no longer function properly. It is also unclear whether it is possible to configure a /32 subnet mask for all operating systems, particularly for IoT devices, industrial machines, and building automation systems. Many convenience features, including the ability to search for services and devices within the same L2 network, will no longer be available.

4.6 OT networks

Operational Technology (OT) networks [27] encompass a broad range of systems, including industrial Internet of Things (IoT), building automation, and control systems. These networks support headless devices that manage critical infrastructure, such as HVAC systems, lighting controls, and automated window shades. Unlike traditional IT networks, which primarily support computing devices with relatively short lifespans, OT networks are designed for long-term stability, often aligning with the lifespan of entire buildings.

One of the primary concerns in OT networking is ensuring maximum stability and reliability. Systems such as alarm systems cannot afford downtime or unpredictable behavior caused by frequent updates. Consequently, the network stack of such OT devices is kept very simple. This is reflected in their minimalistic network configurations, which often rely on static IP addressing rather than dynamic services such as DHCP or DNS, reducing potential points of failure.

Most OT devices operate using IP-based networking, although some legacy systems utilize only Layer 2 protocols. This results in unique challenges when integrating OT with modern IT infrastructure, requiring additional solutions such as gateways to enable Ethernet communication.

Historically, OT networks were completely separate from IT networks to maintain security and reliability, but the increasing need for centralized building management software has driven efforts to merge them.

4.6.1 Challenges in OT networking

One of the primary challenges in OT networks is addressing the presence of «silent» devices that only transmit network packets upon startup and then passively listening. This behavior can create issues in dynamic network environments that rely on DHCP or device tracking databases such as LISP which is part of Cisco's SD-Access solution. If a device does not regularly transmit data, it may be forgotten by the system because the entry is timed out, resulting in communication failures. To mitigate this, administrators employ various workarounds on the devices, such as configuring NTP settings, SNMP traps or setting up dummy syslog destinations to ensure periodic traffic.

In addition, it is very difficult and time-consuming to secure and isolate OT networks well enough to protect against external threats, while ensuring easy management and communication between systems.

Another challenge arises from OT-specific communication protocols, such as BACnet/IP [28], which includes a unique device discovery mechanism using broadcast frames. The introduction of new Software-Defined Systems (for example SD-Access by Cisco) has introduced new considerations, as features such as Layer 2 flooding suppression can disrupt the BACnet/IP's discovery process. To enable communication across subnets, BACnet/IP networks employ BACnet Broadcast Management Devices (BBMDs [29]) that convert broadcast messages into unicast messages for distribution. For this reason, some vendors recommend that networks should not be larger than /24. This is because networks can collapse if too many broadcasts are sent.

Also, some OT devices implement proprietary, slightly different or non-standard variations of common protocols. For example, devices may lack ARP functionality, while some other systems use Ethernet-based communication without full IP support. These inconsistencies complicate integration efforts and may require customized solutions or vendor-specific workarounds.

4.6.2 Future and developments in OT networking

To address security and networking challenges, new standards and technologies are emerging within OT environments. One example is BACnet/SC (Secure Connect) [30], a successor to BACnet/IP that extend the network stack functionality and security. Additionally, BACnet/SC inherently supports DHCP, streamlining device integration and management.

Furthermore, vendors are increasingly integrating REST APIs behind gateways, improving OT device administration and interoperability with IT systems.

5. Important topics in relation to /32 environments

This chapter highlights the resource requirements of the /32 approach for network routers. It also shows how firewalls can be used in a /32 environment and what can be done if large Layer 2 networks are still required.

5.1.1 Hardware limitations related to /32 environment

In order to demonstrate whether a /32 environment approach can be implemented in a campus network with regard to hardware limitations, two different examples are shown below. These simply serve as a brief illustration of how this /32 approach affects the resources of the router and shows roughly what size of networks can be covered. These values are not exact values, but estimates, as the values used may vary depending on the network environment (e.g. number of used uplinks).

The Cisco Catalyst 9300 with 24 available interfaces, which was used in the Cisco test labs, is taken as the reference model for this calculation. The following hardware limitation data can be taken from the router datasheets [31]:

- Number of supported IPv4 routes ➔ 32'000 (24'000 direct routes => A locally connected host prefix and 8'000 indirect routes => A route that is via a remote next hop to reach)
- Number of supported IPv6 routes ➔ 16'000
- Number of supported routing entries in TCAM (used in routers to make routing table lookups very fast) [32], [33] ➔ 8'192
- Number of supported VRF's [34] ➔ 256
- 8 x Catalyst 9300 can be combined in a common stack to increase port density [35]

Example 1

In this example, it is assumed that a total of 3'000 end devices is connected to the Cisco Catalyst 9300 routers. This number puts us in the range of a campus network. Each of these devices need an own routing entry because of the /32 environment approach. For the calculations, it is assumed that the routers are stacked and operated only in IPv4 in each case. This results in a number of available interfaces per stack of 8 x 24 = 192. Every stack has a control plane that corresponds to the hardware limitations of one Catalyst 9300 (5.1.1).

Below is the calculation of the various results:

- Number of required stacks ➔ 3'000 / 192 ≈ 16
- Number of interfaces available ➔ 16 * 192 = 3'072
- Number of IPv4 routes needed per stack ➔ direct routes ≈ 192, indirect routes 15 * 192 ≈ 2'880
- Number of TCAM entries needed per stack ➔ 192 + 2'880 ≈ 3'072

The results show that there are still plenty of reserves. For the values that are growing the most (indirect routes and TCAM), not even half of the possible entries have been reached.

Example 2

This example calculates the maximum number of supported end devices that are possible with the hardware limitation data found of the Catalyst 9300. As in the previous example, stacking and IPv4 is used.

The following maximum values were obtained by testing:

- Max. number of stacks ➔ ≈ 42
- Max. Number of interfaces available ➔ 42 * 192 = 8'064
- Number of IPv4 routes needed ➔ direct routes ≈ 192, indirect routes 41 * 192 ≈ 7'872
- Number of TCAM entries needed ➔ 192 + 7'872 ≈ 8'064

The reason why not more end devices can be connected is due to the number of supported indirect routes (7'872 out of 8'000) and the supported TCAM entries (8'064 out of 8'192), which have almost reached their maximum. However, these results clearly show that many end devices can be operated with a /32 environment. If even more powerful routers are used, the maximum number of end devices can be further increased.

The values will look similar for IPv6, as the TCAM entries will be exhausted first here too. To still be able to connect more end devices, route summarization [36] can be used if this is possible. If, for example, prefix delegation is used in an IPv6 environment, an IPv6 prefix could be assigned for each router. This has the advantage that all hosts connected to this router receive an IPv6 in this prefix. This means that only one IPv6 route needs to be stored on the other routers and not a separate one for each host. Of course, route summarization also works for IPv4.

5.2 Firewall

In a /32 environment, an existing firewall concept can be used without any problems. For example, it is possible to set a separate Anycast Gateway for each department automatically on the corresponding router interface, which is then used to assign an appropriate /32 IP address to the connected host via DHCP. This means that each department still has its own address range, which can be managed with firewall rules.

Of course, it is also possible to use the approach with different VRFs, for example. This allows different networks to be completely separated from each other and managed in separate routing tables. This is desirable, for example, if the network traffic of two customers runs via the same network hardware.

5.3 What if Layer 2 is still needed?

Of course, there will still be some device categories or protocols, such as OT devices (4.6), which only work if they are all operated in the same Layer 2 network. To consider this, alternative solutions must be available for such scenarios that can be integrated into a /32 network. If only a small and manageable number of problem devices are involved, a separate dedicated Layer 2 network can be created for them, which is operated exclusively for this purpose. As a further, but more complex solution, you can also rely on tunnel protocols, like VXLAN, which tunnels Ethernet (Layer 2) traffic over an IP (Layer 3) network. Because VXLAN uses a Layer 3 underlay, there is also no need to deal with the Spanning Tree Protocol (3.6.2).

6. Result discussion

The following chapters summarize, discuss and evaluate the results of the bachelor's thesis.

6.1 Results

By using a virtual and a physical test lab, combined with the products of two different network device manufacturers, it was possible to cover a large test field with many different aspects. It also gave the ability to make targeted and detailed statements about the feasibility of this /32 approach.

The result are two support matrices, which can be found in the Lab Documentation. These matrices provide a quick overview of problems, limitations and restrictions related to /32 environment capabilities. Different colors indicate whether something works without a problem, requires a workaround or is impossible.

The first support matrix deals with the capabilities of the two network device manufacturers Arista and Cisco to set up a /32 environment at all. As can be seen from the matrix, it is currently not possible to implement a functioning and fully usable /32 environment in a campus network. With Arista it fails because not more than one DHCP host can be connected to the same router. With Cisco, the DHCP binding function causes problems, ultimately resulting in chaos in the routing table, as IP addresses of DHCP hosts that are no longer connected remain in the routing table until the DHCP lease has expired.

In addition, the tested enterprise features MAB and 802.1x do not work correctly on the Cisco routers. This showed up with the MAB interface in a way that the Cisco router does not send any Access-Request messages to the RADIUS server and the interface remains in an unauthorized state in any case. On the 802.1x interface a connected MAB host is assigned an IP address by the DHCP server despite the lack of authentication, which does not correspond to the correct behavior. Many of the problems mentioned here are largely due to bugs, which should be able to be resolved with a corresponding fix. Missing functions, such as the automatic setting of static host routes or interface profiles, could be added with appropriate workarounds, for instance with custom shell scripts on third-party systems or on the routers themselves.

The second support matrix shows that the most important protocols (Table 6) in a campus network function correctly in a /32 environment.

In a further step, the same principle of a /32 environment was also simulated in an IPv6 context with a /128 environment. Due to the smaller test setup and the much smaller test scope, the results were not included in a support matrix. However, the initial tests showed that, in contrast to IPv4, a functioning /128 environment can be set up using the help of custom shell scripts. Whether the protocols also work in a /128 environment cannot be confirmed due to a lack of tests.

6.2 Conclusion and outlook

Although the results show that a correctly functioning /32 environment unfortunately cannot be implemented with Arista or Cisco, another conclusion can be drawn: From a purely network-technical point of view, it is possible to set up a /32 environment. Causes preventing this lie exclusively within the router OS and not within any standards, technical limitations or incompatibilities. This could be tested with the various test labs and thus also proven. Investigations also revealed many bugs and

malfunctions in the router OS. This is probably because nobody had comparable requirements beforehand and therefore did not come into such intense contact with the individual functions.

Nevertheless, it has been successfully demonstrated that with this /32 approach, Layer 2, with all its disadvantages, can be successfully reduced to a minimum. This leads to a more stable, more efficient and less error-prone network, which is certainly one of the greatest advantages of a /32 environment. In addition, it could be shown that the Layer 3 routing domain can be successfully extended to the end device.

During the implementation of this project, it was also necessary to learn that third-party systems, such as a DHCP server, must support certain properties in order to successfully distribute a subnet mask of 255.255.255.255 to a DHCP host, for instance. It is therefore important to check in advance whether server services are /32 compatible.

When implementing the /32 environment, workarounds involving custom shell scripts had to be used. Some of these were executed on the DHCP and RADIUS servers, as well as on the routers themselves. These extensions allowed functions to be added to the router OS, or existing ones to be automated. However, such workarounds are not suitable for large campus networks as they are prone to errors and have capacity limitations. In summary, router manufacturers should integrate the necessary functions directly into their systems to ensure secure and reliable operation.

All of the various protocols that have been extensively tested in the /32 environment function properly. Therefore, from a technical protocol point of view, no problems should be expected when implementing a /32 environment. It should be noted, however, that only a small number of the many existing protocols were examined. As many protocols extend beyond Layer 2, this should generally not be an issue, provided they do not depend on Layer 2 multicasts or broadcasts simultaneously. Separate solutions must be set up within a /32 environment for device groups and protocols that are dependent on a common L2 domain.

These two findings on how routers and protocols behave in a /32 environment can be summarized as follows: Once the manufacturers have fixed the discovered bugs and added the missing functions to their OS, significant progress can be made. Fortunately, it seems that the protocols and standards used in the test labs do not need to be adapted, which would be very challenging. Further improvements are also desirable for the RADIUS server in connection with MAB and 802.1x. Currently, the attributes that can be set are exclusively oriented towards Layer 2, e.g. setting a corresponding VLAN. Regarding Layer 3 configurations on interfaces, no attributes can be set. Here too, the manufacturers are in charge to provide suitable vendor-specific attributes to enable VRF and unnumbered interfaces to be set on L3 interfaces, for example.

However, whether these changes and improvements are implemented depends heavily on sufficient demand to make it worthwhile for the manufacturers. This demand certainly also depends on how extensively such a /32 environment, with all its advantages, can be used in other areas of the network. In order to figure this out, the areas of WLAN infrastructures and data centers must be included in future studies. Furthermore, protocol and IPv6 tests would help to create a much clearer scope of what is possible with a /32 environment.

It should be noted that these tests only provide a snapshot of the current situation. In terms of «/32 environment» compatibility the results may differ significantly for other router manufacturers. Nevertheless, this work has provided a comprehensive and detailed overview of the feasibility of a /32 environment.

7. Glossary

Various terms are explained in the glossary below (Table 10).

Technical term	Explanation
BPDU	Control messages used by STP to detect loops and manage topology changes.
Broadcast storm	An overwhelming flood of broadcast traffic that consumes network bandwidth and causes devices to slow down or crash.
BUM	Traffic types in Layer 2 networks that are flooded when destination MAC is not known.
CPE WAN	The WAN-facing interface of a customer device connecting to the provider's network.
Denial-of-service attacks (DoS)	A cyberattack that floods a target system or network with excessive requests, making services unavailable to legitimate users.
ECMP	A routing strategy that uses multiple paths with the same cost to load-balance traffic.
L2 (Layer 2)	The data link layer responsible for direct node-to-node data transfer and MAC addressing.
L3 (Layer 3)	The network layer responsible for IP addressing and routing packets across networks.
Network meltdown	A severe network failure where excessive traffic or cascading faults render the entire network unusable.
OT (Operational Technology)	Technology used to control and monitor industrial equipment and processes.
SPB	A Layer 2 protocol that enables shortest-path forwarding using IS-IS for loop-free topologies.
STP	A protocol that prevents loops in Layer 2 Ethernet networks by creating a loop-free topology.
TCAM	High-speed memory in switches/routers used for fast lookup of ACLs, routing, and QoS rules.
Telnet (alt port)	Popular alternative port to Telnet (for IoT devices)
TRILL	A protocol that combines Layer 2 bridging and Layer 3 routing to optimize Ethernet paths.
TTL	A field in IP packets that limits their lifespan to prevent infinite looping in networks.
VLAN	A logical network that segments Layer 2 traffic on the same physical infrastructure.

VRF	A technology that allows multiple separate routing tables to coexist on a single router.
VXLAN	An overlay protocol that encapsulates Layer 2 frames in UDP packets to span L2 across L3 networks.
WLAN	Wireless Local Area Network allows devices to connect to a network using radio waves.
WSD	A Microsoft protocol used to automatically discover network devices and services.

Table 10: Glossary

8. List of helper tools

Table 11 lists all the tools that were used for the bachelor's thesis.

Task area	Tools
Literature research and management	Google, ChatGPT
Idea generation	ChatGPT
Translations	DeepL, ChatGPT
Text creation, text optimization, spelling and grammar check	Word, DeepL, ChatGPT

Table 11: List of helper tools

9. List of illustrations

10. List of tables

11. List of references

[1] “Spanning Tree (STP) Limitations.” Accessed: Apr. 22, 2025. [Online]. Available: https://networklessons.com/spanning-tree/spanning-tree-stp-limitations

[2] I. Pepelnjak, “Transparent Bridging (aka L2 Switching) Scalability Issues.” Accessed: Mar. 12, 2025. [Online]. Available: https://blog.ipspace.net/2012/05/transparent-bridging-aka-l2-switching/

[3] I. Pepelnjak, “STP loops strike again « ipSpace.net blog.” Accessed: Mar. 12, 2025. [Online]. Available: https://blog.ipspace.net/2012/04/stp-loops-strike-again/

[4] I. Pepelnjak, “Layer-2 Network Is a Single Failure Domain « ipSpace.net blog.” Accessed: Mar. 12, 2025. [Online]. Available: https://blog.ipspace.net/2012/05/layer-2-network-is-single-failure/

[5] I. Pepelnjak, “Transparent Bridging (aka L2 Switching) Scalability Issues « ipSpace.net blog.” Accessed: Mar. 12, 2025. [Online]. Available: https://blog.ipspace.net/2012/05/transparent-bridging-aka-l2-switching/

[6] L. Bell, A. Smith, P. Langille, A. Rijhsinghani, and K. McCloghrie, “Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions.” Accessed: Apr. 22, 2025. [Online]. Available: https://www.ietf.org/rfc/rfc2674.txt

[7] S. Luber, “Was ist TRILL (Transparent Interconnection of Lots of Links)?” Accessed: Mar. 12, 2025. [Online]. Available: https://www.ip-insider.de/was-ist-trill-transparent-interconnection-of-lots-of-links-a-f917354a1857046bbfbbc2f7a4085528/

[8] Tom, “The Death of TRILL | The Networking Nerd.” Accessed: Mar. 12, 2025. [Online]. Available: https://networkingnerd.net/2016/05/11/the-death-of-trill/

[9] “Software-Defined Access Solution Design Guide - Cisco.” Accessed: Mar. 06, 2025. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-design-guide.html#Underlaynetwork

[10] R. Molenaar, “Cisco Campus Network Design Basics.” Accessed: Apr. 21, 2025. [Online]. Available: https://networklessons.com/cisco/ccna-routing-switching-icnd1-100-105/cisco-campus-network-design-basics

[11] A. Viescinski, “Network Traffic: North-South and East-West | Baeldung on Computer Science.” Accessed: Apr. 26, 2025. [Online]. Available: https://www.baeldung.com/cs/network-traffic-north-south-east-west

[12] M. Hossain, “Trends in Data Center Security: Part 1 – Traffic Trends - Cisco Blogs.” Accessed: Apr. 26, 2025. [Online]. Available: https://blogs.cisco.com/security/trends-in-data-center-security-part-1-traffic-trends

[13] “Network Traffic Shifts from East/West to North/South | Imagit Inc.” Accessed: Apr. 26, 2025. [Online]. Available: https://www.imagit.com/network-traffic-shifts-from-east-west-to-north-south/

[14] A. H. Villa and E. Varki, “Characterization of a campus Internet workload,” 2011.

[15] M. A. Mohd Ariffin, “Network Traffic Profiling Using Data Mining Technique in Campus Environment,” *International Journal of Advanced Trends in Computer Science and Engineering*, vol. 9, no. 1.3, pp. 422–428, Jun. 2020, doi: 10.30534/ijatcse/2020/6691.32020.

[16] M. Haffey, “Characterization of Periodic Network Traffic,” 2017, doi: 10.11575/PRISM/25285.

[17] S. Aggarwal, “Top 40 Protocols: A Comprehensive Guide.” Accessed: Apr. 26, 2025. [Online]. Available: https://blog.certcube.com/top-40-protocols-a-comprehensive-guide/

[18] N. Greene, “16 Most Common Network Protocols You Should Know.” Accessed: Apr. 26, 2025. [Online]. Available: https://www.auvik.com/franklyit/blog/common-network-protocols/

[19] K. Yasar and M. Goss, “15 Common Network Protocols and Their Functions Explained.” Accessed: Apr. 26, 2025. [Online]. Available: https://www.techtarget.com/searchnetworking/feature/12-common-network-protocols-and-their-functions-explained

[20] R. Molenaar, “ARP (Address Resolution Protocol) explained.” Accessed: Feb. 26, 2025. [Online]. Available: https://networklessons.com/ip-services/arp-address-resolution-protocol-explained

[21] “ATA over Ethernet - Wikipedia.” Accessed: May 25, 2025. [Online]. Available: https://en.wikipedia.org/wiki/ATA_over_Ethernet

[22] “Fibre Channel over Ethernet - Wikipedia.” Accessed: May 25, 2025. [Online]. Available: https://en.wikipedia.org/wiki/Fibre_Channel_over_Ethernet

[23] “Phillips Hue Bridge over VLAN - Fortinet Community.” Accessed: Mar. 01, 2025. [Online]. Available: https://community.fortinet.com/t5/Support-Forum/Phillips-Hue-Bridge-over-VLAN/td-p/72183

[24] “Apple Bonjour und Einschränkungen für mehrere VLANs auf Dell Networking-Switches | Dell Schweiz.” Accessed: Feb. 26, 2025. [Online]. Available: https://www.dell.com/support/kbdoc/de-ch/000142540/apple-bonjour-und-einschr%C3%A4nkungen-f%C3%BCr-mehrere-vlans-auf-dell-networking-switches

[25] dbellamk, “Cisco SD-Access Layer2 flooding - Cisco Community.” Accessed: Mar. 01, 2025. [Online]. Available: https://community.cisco.com/t5/networking-knowledge-base/cisco-sd-access-layer2-flooding/ta-p/3943916

[26] R. Parmar, “Joining Process of an Cisco Access Point - Cisco Community.” Accessed: Mar. 01, 2025. [Online]. Available: https://community.cisco.com/t5/wireless-mobility-knowledge-base/joining-process-of-an-cisco-access-point/ta-p/3149279

[27] E. Banks, D. Conry-Murray, and G. Ferro, “Managing OT Networks.” Accessed: Mar. 05, 2025. [Online]. Available: https://packetpushers.net/podcasts/heavy-networking/hn735-managing-ot-networks/

[28] “BACnet/IP - Feldbusanschlüsse (Konnektivität) | Feldbusanschlüsse | ABB.” Accessed: Mar. 05, 2025. [Online]. Available: https://new.abb.com/drives/de/konnektivitaet/feldbusanschluesse/bacnet-ip

[29] “BACnet Broadcast Management Device.” Accessed: Mar. 05, 2025. [Online]. Available: https://infosys.beckhoff.com/index.php?content=../content/1031/tcbacnet/12748445963.html&id=

[30] “Die Erweiterung des BACnet-Standards bei SAUTER.” Accessed: Mar. 05, 2025. [Online]. Available: https://www.sauter-building-control.ch/innovation/bacnet-sc/

[31] “Catalyst 9300 Series Switches Data Sheet - Cisco.” Accessed: Jun. 10, 2025. [Online]. Available: https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9300-series-switches/nb-06-cat9300-ser-data-sheet-cte-en.html

[32] A. Taylor, A. A. Ferguson, and K. Gabrielsen, “Understand IPv4 Hardware Resources on Catalyst 9000 Switches - Cisco.” Accessed: Jun. 10, 2025. [Online]. Available: https://www.cisco.com/c/en/us/support/docs/switches/catalyst-9300-series-switches/217714-understand-ipv4-hardware-resources-on-ca.html#toc-hId--267947411

[33] “Memory - TCAM Lookups.” Accessed: Jun. 10, 2025. [Online]. Available: https://notes.networklessons.com/memory-tcam-lookups

[34] “Products - Cisco Catalyst 9000 Switching Platform FAQ - Cisco.” Accessed: Jun. 10, 2025. [Online]. Available: https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/nb-06-cat9k-swit-plat-faq-cte-en.html

[35] “Catalyst 9300 Stackwise System Architecture White Paper - Cisco.” Accessed: Jun. 10, 2025. [Online]. Available: https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9300-series-switches/white-paper-c11-741468.html

[36] R. Molenaar, “Introduction to Route Summarization.” Accessed: Jun. 10, 2025. [Online]. Available: https://networklessons.com/rip/introduction-route-summarization

Ostschweizer Fachhochschule, Standort Rapperswil