Enterprise Secure Network Design and Implementation: A VLAN-Segmented, ACL-Controlled Campus Network for a Computer Science and Engineering Department

Suraj Jalindar Bedekar[1], Vrushabh Gangadhar Zingade[2], Swapnil Gajanan Ghuge[3], Swapnil Prakash Gandhe[4] Prof. S. S. Mule[5]
[1,2,3,4]*Department of Computer Science & Engineering, MIT College of Railway Engineering and Research (MITCORER), Barshi, Maharashtra*
[5]*Project Guide, Department of Computer Science & Engineering, MIT College of Railway Engineering and Research (MITCORER), Barshi, Maharashtra*

***Abstract*—Flat departmental networks create large broadcast domains, offer no internal segmentation and become hard to manage as they grow. This paper presents the design, implementation and verification, in Cisco Packet Tracer, of an Enterprise Secure Network for a Computer Science and Engineering (CSE) department, based on a hierarchical, policy-controlled architecture rather than any new protocol. Eleven laboratory VLANs (10 to 20) in three building zones reach a Cisco 3560 Layer-3 core switch over 802.1Q trunks, which provides hardware inter-VLAN routing through SVI gateways, DHCP relay and a routed transit link to a Cisco 2911 router. A centralised Server Farm (VLAN 30) hosts DNS, Syslog, TFTP, DHCP and FTP services, and a protected Faculty and Admin VLAN (VLAN 90) has full access. The extended ACL SECURE-LABS blocks students from the admin network and from the Syslog and TFTP servers while keeping DNS, DHCP and FTP available. Sixteen test cases covering connectivity, services and access control were defined, and an analytical comparison with a flat design show that the average number of hosts receiving each broadcast frame falls from 267 to 23.4, a reduction of about 91%. The results show that segmentation, central services and least-privilege policy can be combined in one verified, reproducible design that can later be extended with redundancy, OSPF, NAT and Layer-2 security.**

***Index Terms*—Enterprise network, VLAN, inter-VLAN routing, DHCP relay, access control list, server farm, Packet Tracer, network security.**

I. INTRODUCTION

A campus network is the backbone of every modern academic institution. In a CSE department it must connect practical laboratories, a computer centre, faculty and administrative staff, printers and wireless users, and give all of them dependable access to shared services such as address assignment, name resolution, logging, configuration backup and file transfer. As users and devices multiply, the way the network is organised decides how secure, fast and manageable it remains.

Early departmental networks are typically flat: all computers share one or two large Layer-2 segments behind a single router. Every broadcast then reaches every host, any computer can attempt to reach any server, and inter-VLAN traffic must pass through an external router. Vendor guidance therefore recommends a structured, hierarchical design in which access switches serve end users, a Layer-3 core performs routing and traffic is segmented by function [1], [2].

Several mature technologies make this practical: IEEE 802.1Q VLANs divide a switched network into logical segments [3]; private IPv4 addressing gives a structured address plan [4]; DHCP with a relay agent delivers addresses from a central server [5]; Syslog and TFTP provide central logging and configuration backup [6], [7]; and Access Control Lists (ACLs) decide which communication is permitted. This work applies them to one realistic scenario, a CSE department of about 270 laboratory computers, and

implements and tests the complete design in Cisco Packet Tracer [8].

The key objectives of this research are:

- To design a multi-zone network covering the West Wing, Computer Centre and East Wing laboratories, a Server Farm and an administrative zone.
- To give each of the eleven laboratories its own VLAN and /24 subnet, joined to the core by 802.1Q trunks.
- To implement hardware inter-VLAN routing with SVIs on a Cisco 3560 Layer-3 core and a routed transit link to Router1.
- To provide centralised DNS, DHCP (with relay), Syslog, TFTP and FTP services in VLAN 30.
- To enforce a security policy with an extended ACL that gives VLAN 90 full access and blocks students from the admin network and from the Syslog and TFTP servers.
- To verify connectivity, services and policy through documented tests and to compare the design quantitatively with a flat network.

II. RELATED WORK

Foundational concepts of switching, routing and network security are drawn from standard textbooks [9], [10]. Cisco's campus design guidance and the CCNA curriculum describe a hierarchical access, distribution and core model that can be collapsed into two tiers for medium-sized sites [1], [2]. IEEE 802.1Q defines VLAN tagging and bridging [3], but VLANs alone enforce no policy: inter-VLAN traffic is controlled only when a Layer-3 device applies ACLs, and misconfigured trunks can allow VLAN hopping. Rapid Spanning Tree [11] and link aggregation [12] give fast convergence and extra bandwidth, and OSPFv2 [13] is the standard interior routing protocol for campuses, while IEEE 802.1X admits only authenticated devices but needs an AAA infrastructure [14].

At the policy level, NIST SP 800-41 recommends default-deny rule sets placed at network boundaries and between internal segments [15], and commercial next-generation firewalls add inspection and intrusion prevention at significant cost [16]. Intrusion detection and prevention systems add detection capability but need tuning [17]. Zero-trust architecture [18] treats no network location as implicitly trusted, and micro-segmentation with default-deny rules is one of its practical building blocks; this project adopts that idea in a limited form. Software-defined networking [19] offers central policy but makes the controller a critical component. Supporting services such as DHCP [5], Syslog [6], TFTP [7], SNMP [20], NTP [21] and monitoring platforms [22] are mostly unencrypted by default and must be confined to a protected segment. Wireless controllers using CAPWAP [23] and virtualised infrastructure services [24] centralise management but add dependencies. Table 1 summarises these works.

Table 1 Literature Review Summary

Ref.	Technique	Limitation
[1], [2]	Hierarchical campus design	Vendor-centric; fragments, not a full case
[3]	IEEE 802.1Q VLAN tagging	No policy without L3 ACLs; VLAN hopping
[11], [12]	RSTP, LACP	Does not protect against device failure
[13]	OSPFv2 link-state routing	Area and cost design needed
[14]	IEEE 802.1X, RADIUS	Needs AAA; IoT lacks supplicants
[15], [16]	Default-deny firewalls, NGFW	Cost; perimeter misses lateral traffic
[17]	Signature / anomaly IDPS	False positives; tuning
[18]	Zero-trust architecture	Needs identity-aware policy engine
[19]	Software-defined networking	Controller risk; migration cost
[5]–[7], [20]–[22]	DHCP, Syslog, TFTP, SNMP, NTP, NMS	Mostly unencrypted; must be isolated
[23]	CAPWAP wireless control	Controller dependency
[24]	Hypervisor consolidation	Host failure affects many services

Every technology a secure campus network needs is therefore mature. The gap lies in how they are brought together and verified: guides present protocols in isolation, student exercises usually stop at VLANs and inter-VLAN routing, and published designs seldom include a test plan that proves each policy and a quantitative comparison with the flat design they replace. This work addresses that gap; its contribution is integration and verification, not a new protocol.

III. SYSTEM ARCHITECTURE

The proposed network replaces the flat, router-centric model with a hierarchical, segmented and policy-controlled one. It follows four principles: separate traffic by function; route at the core in hardware; keep shared services in one protected Server Farm; and permit only the traffic each group needs. As Fig. 1 shows, eleven Cisco 2960 access switches in three zones connect over 802.1Q trunks to a Cisco 3560 Layer-3 core. A routed /30 transit link joins the core to Router1 (Cisco 2911), which is the gateway of the Server Farm switch carrying VLAN 30 and VLAN 90.

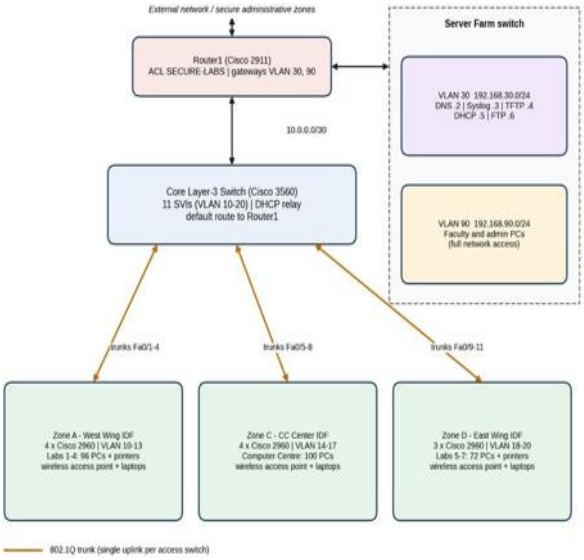


Fig. 1 Architecture of the implemented CSE Department network

Subnets use private addresses [4] of the form 192.168.<VLAN>.0/24 with the gateway at .1. Table 2 gives the allocation. With h host bits a subnet holds

$$H = 2^{*h*} - 2 \quad (1)$$

hosts, so a /24 holds 254. A lab of 24 hosts uses 9.4% of its subnet, and across the eleven lab subnets 268 hosts use 9.6% of the 2,794 usable addresses. A /27 would be tighter, but /24 subnets were chosen for a uniform plan, simple SVIs and room for growth.

Table 2 Vlan and Ip Allocation

Zone	VLANs	Subnets	Hosts
West Wing labs	10–13	192.168.10–13.0/24	96
Computer Centre	14–17	192.168.14–17.0/24	100
East Wing labs	18–20	192.168.18–20.0/24	72
Server Farm	30	192.168.30.0/24	5 servers
Faculty / Admin	90	192.168.90.0/24	3 PCs
Core–Router1 transit	–	10.0.0.0/30	.1 core, .2 Router1

IV. SYSTEM MODULES

The implementation is divided into five modules, built in this order so that each layer could be tested before the next was added.

4.1 Access Layer

Each of the eleven lab switches carries a single VLAN on its access ports (Fa0/1–23) and one 802.1Q trunk (Fa0/24) to the core, and forwards its logs to the Syslog server. Only the VLAN number and name differ between switches.

4.2 Layer-3 Core

On the Cisco 3560 core, ip routing is enabled and an SVI is created per lab as its default gateway. The command ip helper-address 192.168.30.5 on every SVI converts the clients' DHCP broadcasts into unicasts for the central DHCP server. The port facing Router1 is a routed port (no switchport) in 10.0.0.0/30, and a default route of last resort points to Router1.

```
interface Vlan10
 ip address 192.168.10.1 255.255.255.0
 ip helper-address 192.168.30.5
interface GigabitEthernet0/1
 no switchport
 ip address 10.0.0.1 255.255.255.252
ip route 0.0.0.0 0.0.0.0 10.0.0.2
```

4.3 Router1 and the SECURE-LABS ACL

Router1 sits at the edge of the department network and is also the gateway of VLAN 30 and VLAN 90. Because an ACL is read from top to bottom and stops at the first match, the rules are ordered so that (1) VLAN 90 receives unrestricted access; (2) established TCP and ICMP echo-replies are allowed, so replies to admin-initiated traffic return; (3) students are denied the admin subnet; (4) students are denied the Syslog server 192.168.30.3 and the TFTP server 192.168.30.4; and (5) a final explicit permit keeps DNS, DHCP and FTP working. Table 3 shows the resulting policy matrix.

```
ip access-list extended SECURE-LABS
 permit ip 192.168.90.0 0.0.0.255 any
 permit tcp any any established
 permit icmp any any echo-reply
 deny ip any 192.168.90.0 0.0.0.255
 deny ip any host 192.168.30.3
 deny ip any host 192.168.30.4
 permit ip any any
```

Table 3 Access Control Policy Matrix

From \ To	VLAN 90	DNS, DHCP, FTP	Syslog, TFTP	Other labs
Students (VLAN 10–20)	Deny	Permit	Deny	Permit
Faculty/admin (VLAN 90)	–	Permit	Permit	Permit

Traffic between lab VLANs is routed directly by the core and is not filtered; all other traffic leaves the core through Router1, where the ACL applies.

4.4 Server Farm

Five Server-PT devices on a dedicated switch in VLAN 30 provide the services in Table 4. The DNS server resolves names such as ftp.project.local to 192.168.30.6 and its address is embedded in every DHCP pool. The Syslog server (UDP 514) collects events from all Cisco devices [6]. The TFTP server stores configuration backups taken with copy running-config tftp [7]. The DHCP server runs one pool per VLAN, for example 192.168.10.100 onward for VLAN 10, leaving the lower addresses free for static devices [5]. The FTP server provides authenticated file sharing over TCP.

Table 4 Server Farm Address Plan

Server	Service	IPv4 address (gateway .1)
Server0	DNS	192.168.30.2/24
Server1	Syslog	192.168.30.3/24
Server2	TFTP	192.168.30.4/24
Server3	DHCP	192.168.30.5/24
Server4	FTP	192.168.30.6/24

4.5 Packet Decision and DHCP Relay

Fig. 2 follows a packet from a lab PC. The core routes traffic between lab VLANs directly; everything else is forwarded to Router1, where the ACL decides. Fig. 3 shows address assignment: the client's broadcast reaches the core SVI, which relays it as a unicast through Router1 to the DHCP server in VLAN 30 [5].

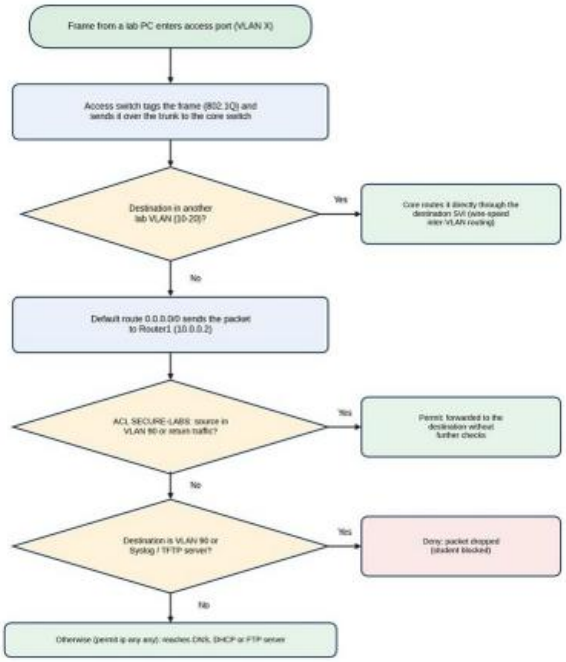


Fig. 2 Packet processing and policy enforcement

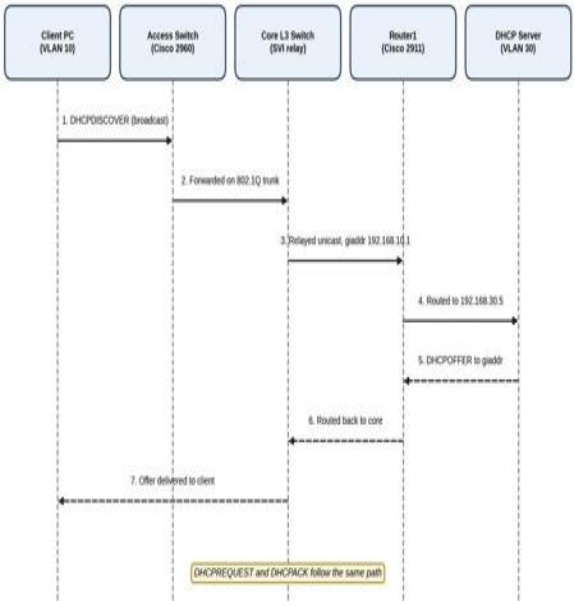


Fig. 3 DHCP exchange through the relay agent

4.6 Key Features

- One VLAN and /24 subnet per laboratory, with hardware inter-VLAN routing at the core
- Central DHCP with relay, DNS, Syslog, TFTP and FTP in a protected Server Farm
- Least-privilege ACL protecting the admin network and infrastructure servers
- Spare /24 capacity: a new lab needs only one

VLAN, one SVI and one DHCP pool
- Per-VLAN fault isolation and central logs for faster troubleshooting

V. METHODOLOGY

The project followed a hybrid methodology that combines structured, waterfall-style planning with iterative refinement driven by test results:
- Requirement analysis: identifying the limitations of the flat network and writing nine functional and five non-functional requirements (security, performance, scalability, manageability, maintainability).
- System design: the two-tier architecture, VLAN and IP plan, device roles, policy matrix and packet-decision algorithm.
- Implementation in Cisco Packet Tracer 8.x [8]: the topology was built in six steps: place and cable devices; create VLANs and trunks; enable routing, SVIs, relay and the transit link; configure Router1 and the ACL; configure the five servers; enable logging and back up configurations.
- Testing: sixteen test cases (Table 5) using ping, nslookup, ftp and show commands from representative hosts in every VLAN, three VLAN 90 PCs and the five servers.
- Analysis: analytical comparison with a flat design for broadcast exposure and uplink oversubscription.

Table 5 Test Cases

ID	Test	Expected result
TC-01	show vlan brief on a lab switch	Lab VLAN present, ports assigned
TC-02	show ip route on the core	11 connected subnets, default route
TC-03/04	DHCP on a PC in VLAN 10 and in VLANs 11–20	Lease from own pool
TC-05	VLAN 10 PC pings VLAN 12 PC	Success via core SVIs
TC-06	nslookup ftp.project.local	Returns 192.168.30.6
TC-07	FTP login as admin	File directory accessible
TC-08	Shut / enable a core interface	Messages logged on Syslog
TC-09	Backup config to TFTP	File appears on server
TC-10	VLAN 90 pings labs and servers	All succeed
TC-11	Lab PC pings a VLAN 90 PC	Blocked by ACL
TC-12/13	Lab PC pings Syslog / TFTP server	Blocked
TC-14	Lab PC uses DNS, DHCP, FTP	Services work
TC-15	show access-lists after tests	Hit counters rise
TC-16	Wireless laptop joins zone AP	Gets address, reaches servers

V. RESULTS AND DISCUSSION

6.1 Functional Evaluation

The test cases of Table 5 were executed in Packet Tracer. The core routing table confirmed that all eleven SVIs were
active and that the default route of last resort pointed to Router1 (TC-02). Lab PCs obtained leases, gateways and the DNS address from the central DHCP server through the relay; names resolved through the central DNS server; FTP, TFTP and Syslog operated as designed; and the SECURE-LABS ACL behaved as specified, with hit counters rising on its deny and permit lines (TC-15, Fig. 5).

6.2 Broadcast Exposure

In a flat network of N = 268 lab hosts, each broadcast reaches N − 1 = 267 hosts. With VLANs of n_i hosts, the average number of receivers is

$$E_{seg} = \Sigma\, n_i(n_i - 1) / N \quad (2)$$

For four VLANs of 24 hosts, four of 25 and three of 24, $E_{seg} = 6264 / 268 \approx 23.4$, a reduction of $1 - 23.4/267 \approx 91\%$
(Fig. 4, Table 6).

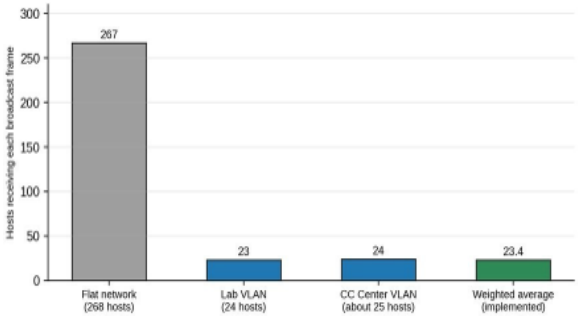


Fig. 4 Hosts receiving each broadcast: flat network versus VLANs

6.3 Uplink Oversubscription

For an access switch with P ports of speed B_p sharing U uplinks of speed B_u,

$$R = (P \times B_p) / (U \times B_u) \quad (3)$$

A 24-port FastEthernet lab switch with one 100 Mbps trunk gives R = 24:1 when all ports are busy, slightly above the common rule of thumb of about 20:1. A Gigabit uplink would reduce it to 2.4:1 and is therefore a recommended improvement.

Table 6 Flat Versus Implemented Network

Metric	Flat design	Implemented design
Broadcast domains	1 of 268 hosts	11 VLANs of 24–25 hosts
Avg. broadcast receivers	267	23.4 (about 91% fewer)
Inter-VLAN routing	External router	L3 core, hardware routing
Student access to admin	Open	Blocked by ACL
Student access to Syslog, TFTP	Open	Blocked
Addressing, names	Manual, per host	Central DHCP and DNS
Logging, backup	Per device	Central Syslog, TFTP

6.4 Test Evidence

Fig. 5 shows two captures from the test bed: a lab PC reaching the DHCP server through the ACL (TC-14) and the SECURE-LABS access list on Router1 with its match counters (TC-15).

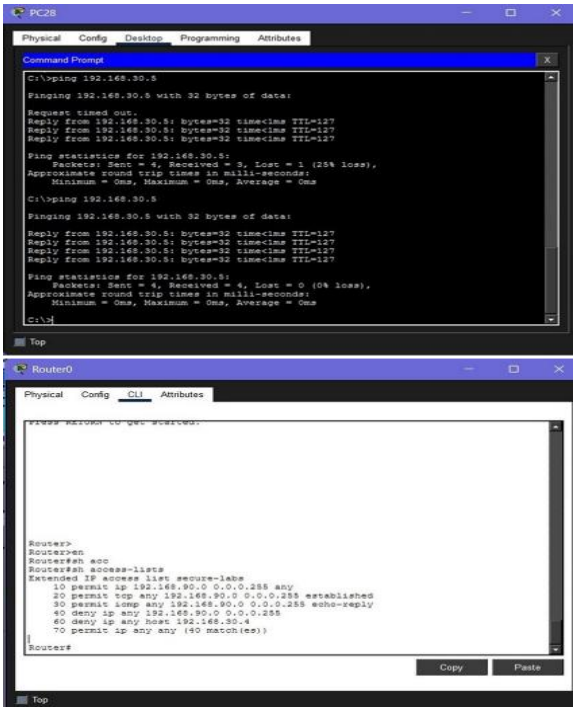


Fig. 5 (a) Lab PC reaching the DHCP server (TC-14); (b) SECURE-LABS ACL on Router1 (TC-15)

6.5 Discussion

The analysis shows that performance, manageability and security can be achieved together by integrating standard technologies. Two weaknesses were also found and are recorded as future work: traffic between lab VLANs is not filtered because the core routes it directly, and each access switch has a single 100 Mbps uplink. Further limitations are that all results come from simulation, so real-hardware throughput and latency under load were not measured; the ACL is stateless; there is no device-level redundancy; routing is static; and FTP and TFTP are unencrypted, with demonstration-only accounts. Table 7 maps each objective to its outcome.

Table 7 Objectives and Achieved Outcomes

Objective	Outcome
Multi-zone design	Three zones, Server Farm and admin VLAN
VLAN per laboratory	11 VLANs (10–20), one /24 each, 802.1Q trunks
Inter-VLAN routing	11 SVIs on the 3560 core; /30 transit to Router1
Central services	DNS, DHCP + relay, Syslog, TFTP, FTP in VLAN 30
Security policy	SECURE-LABS: admin and Syslog/TFTP protected
Verification	16 test cases; broadcast exposure down about 91%

VII. CONCLUSION AND FUTURE WORK

This paper presented a segmented, policy-controlled network for a CSE department, implemented and tested in Cisco Packet Tracer. Eleven laboratory VLANs are routed by a Cisco 3560 Layer-3 core, a centralised Server Farm provides DNS, DHCP, Syslog, TFTP and FTP, and Router1 enforces the SECURE-LABS ACL that gives VLAN 90 full access while blocking students from the admin network and from the Syslog and TFTP servers. Segmentation reduces broadcast exposure by about 91% compared with a flat network. The design is a template for college and small-enterprise campuses, a teaching reference for VLAN, DHCP relay and ACL laboratories, and a baseline for planning a real deployment.

Future work includes Gigabit uplinks with EtherChannel and a second core switch with HSRP or VRRP; dynamic routing with OSPF, NAT at the edge

and a next-generation firewall with IDS/IPS; Layer-2 security (port security, DHCP snooping, dynamic ARP inspection and 802.1X); ACLs between lab VLANs, SSH-only management, SNMP and NTP; replacing FTP and TFTP with SFTP and adding centralised AAA; and automation with Python or Ansible together with validation on real Cisco equipment.

ACKNOWLEDGMENT

The authors thank the Department of Computer Science & Engineering, MIT College of Railway Engineering and Research, Barshi, for its guidance and support, and the Cisco Networking Academy for the Packet Tracer simulator used in this work.

REFERENCES

[1] Cisco Systems, "Cisco validated design: Campus LAN and WLAN design guide," [Online]. Available: https://www.cisco.com/go/designzone.

[2] W. Odom, *CCNA 200-301 Official Cert Guide*, vol. 1. Indianapolis, IN, USA: Cisco Press, 2019.

[3] *IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks*, IEEE Std. 802.1Q-2018, Jul. 2018.

[4] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address allocation for private internets," IETF, RFC 1918, Feb. 1996.

[5] R. Droms, "Dynamic host configuration protocol," IETF, RFC 2131, Mar. 1997.

[6] R. Gerhards, "The syslog protocol," IETF, RFC 5424, Mar. 2009.

[7] K. Sollins, "The TFTP protocol (revision 2)," IETF, RFC 1350, Jul. 1992.

[8] Cisco Networking Academy, "Cisco Packet Tracer," [Online]. Available: https://www.netacad.com/cisco-packet-tracer.

[9] S. Tanenbaum and D. J. Wetherall, *Computer Networks*, 5th ed. Boston, MA, USA: Pearson, 2011.

[10] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach*, 8th ed. Hoboken, NJ, USA: Pearson, 2021.

[11] *IEEE Standard for Local and Metropolitan Area Networks—Media Access Control (MAC) Bridges*, IEEE Std. 802.1D-2004, Jun. 2004.

[12] *IEEE Standard for Local and Metropolitan Area Networks—Link Aggregation*, IEEE Std. 802.1AX-2020, 2020.

[13] J. Moy, "OSPF version 2," IETF, RFC 2328, Apr. 1998.

[14] *IEEE Standard for Local and Metropolitan Area Networks—Port-Based Network Access Control*, IEEE Std. 802.1X-2020, Feb. 2020.

[15] K. Scarfone and P. Hoffman, "Guidelines on firewalls and firewall policy," NIST, Gaithersburg, MD, USA, SP 800-41 Rev. 1, Sep. 2009.

[16] Fortinet, Inc., "FortiGate next-generation firewall," [Online]. Available: https://www.fortinet.com/products/next-generation-firewall.

[17] K. Scarfone and P. Mell, "Guide to intrusion detection and prevention systems (IDPS)," NIST, Gaithersburg, MD, USA, SP 800-94, Feb. 2007.

[18] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," NIST, Gaithersburg, MD, USA, SP 800-207, Aug. 2020.

[19] D. Kreutz, F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proc. IEEE*, vol. 103, no. 1, pp. 14–76, Jan. 2015.

[20] D. Harrington, R. Presuhn, and B. Wijnen, "An architecture for describing SNMP management frameworks," IETF, RFC 3411, Dec. 2002.

[21] D. Mills, J. Martin, J. Burbank, and W. Kasch, "Network time protocol version 4: Protocol and algorithms specification," IETF, RFC 5905, Jun. 2010.

[22] Paessler GmbH, "PRTG Network Monitor," [Online]. Available: https://www.paessler.com/prtg.

[23] P. Calhoun, M. Montemurro, and D. Stanley, Eds., "Control and provisioning of wireless access points (CAPWAP) protocol specification," IETF, RFC 5415, Mar. 2009.

[24] VMware by Broadcom, "VMware vSphere," [Online]. Available: https://www.vmware.com/products/cloud-infrastructure/vsphere.