Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Experimental Implementation of Network Access Control with Dynamic VLAN Assignment Using Active Directory, Windows NPS, RADIUS, and IEEE 802.1X

Abdulraheem Hamoud Mohammed Almaweri

abdulraheemalmaweri@gmail.com

Académie des sciences et technologies industrielles (ASTI)

Research Article

Keywords: Network Access Control, IEEE 802.1X, RADIUS, Network Policy Server, Active Directory, Dynamic VLAN, PEAP, EAP-MSCHAPv2, AD CS, Cisco IOSvL2.

Posted Date: September 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11107426/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Experimental Implementation of Network Access Control with Dynamic VLAN Assignment Using Active Directory, Windows NPS, RADIUS, and IEEE 802.1X

Abdulraheem Hamoud Mohammed Almaweri
Académie des sciences et technologies industrielles (ASTI), Department of Computer Science and Networking/Cybersecurity
Next to the Airport, Bicol, Djibouti, Republic of Djibouti
Corresponding author: Abdulraheem Hamoud Mohammed Almaweri
E-mail: abdulraheemalmaweri@gmail.com
ORCID: https://orcid.org/0009-0004-5386-3800

Abstract

Enterprise networks increasingly require access-control mechanisms that authenticate users before permitting access and dynamically enforce network segmentation according to user identity and organisational role. This study presents an end-to-end implementation and empirical evaluation of a Network Access Control (NAC) architecture integrating Active Directory Domain Services (AD DS), Microsoft Network Policy Server (NPS), RADIUS, IEEE 802.1X, Protected Extensible Authentication Protocol (PEAP) with EAP-MSCHAPv2, Active Directory Certificate Services (AD CS), and Cisco IOSvL2 switching. The architecture maps Active Directory security-group membership to dynamically assigned VLANs through RADIUS authorisation attributes. Two user groups, VLAN10-Users and VLAN20-Users, are mapped to VLAN 10 and VLAN 20, respectively. A quarantine VLAN, VLAN 999, is additionally configured to contain authentication failures and authentication-server unavailability. The implementation was evaluated in a virtualised VMware/GNS3 laboratory environment consisting of Windows Server, Windows 10, Cisco IOSvL2, and Cisco 3725 components. Functional scenarios included successful authentication, invalid credentials, unauthorised users, dynamic VLAN assignment, and NPS/RADIUS server failure. Across 25 authentication events spanning five scenarios, valid authentication achieved a 100% observed success rate. Mean authentication latency was 1.42 s for VLAN 10 (observed range 1.25–1.68 s) and 1.48 s for VLAN 20 (1.28–1.74 s). Invalid credential attempts were rejected with a 100% observed rejection rate, with no observed access to production VLANs. During the NPS/RADIUS failure scenario, the configured dead-server mechanism placed the endpoint into VLAN 999 within 4.82 s under the experimental conditions. No unauthorised production-VLAN access was observed during the failure tests. The results provide experimental evidence that Active Directory group membership can be translated into network-level authorisation through NPS and RADIUS and enforced dynamically by a Cisco IOSvL2 authenticator using IEEE 802.1X. The study provides a reproducible laboratory model for cybersecurity education, network-access-control experimentation, and identity-based network segmentation research.

Keywords:

Network Access Control; IEEE 802.1X; RADIUS; Network Policy Server; Active Directory; Dynamic VLAN; PEAP; EAP-MSCHAPv2; AD CS; Cisco IOSvL2.

1. Introduction

Enterprise networks carry traffic belonging to populations of users whose access requirements differ substantially, and the boundaries between those populations are increasingly expected to be enforced by the network itself rather than by the applications that run over it. Conventional Layer 2 segmentation associates each switch port with a virtual local area network (VLAN) that is decided in advance and written into the port configuration. A VLAN allocated in this way describes where a device is plugged

in; it says nothing about who is using it. Any person with physical access to an active port therefore inherits whatever authorisation that port carries, and the segmentation that the VLAN was intended to provide is only as strong as the physical security of the wiring closet.

IEEE 802.1X addresses the first half of this problem by defining port-based network access control, in which a switch port remains unauthorised until a supplicant has been authenticated through an authentication server [1]. The Extensible Authentication Protocol (EAP) carries the authentication exchange [2], the Remote Authentication Dial In User Service (RADIUS) protocol transports it between the authenticator and the authentication server [3], and RFC 3579 specifies how EAP is encapsulated within RADIUS [4]. In a Windows environment the authentication server role is commonly filled by Network Policy Server (NPS), which evaluates requests against policies expressed in terms of Active Directory objects [5], [6].

Authentication on its own, however, only decides whether a port is opened. It does not decide what the port is opened onto. The second half of the problem is authorisation: translating a verified identity into a specific network position. RFC 2868 supplies the mechanism by defining the Tunnel-Type, Tunnel-Medium-Type and Tunnel-Private-Group-ID attributes, which allow a RADIUS server to return a VLAN identifier alongside an Access-Accept message [7]. When an authenticator applies those attributes, the VLAN follows the authenticated user rather than the physical port, and the VLAN segmentation becomes identity-driven. This is the mechanism on which the architecture evaluated in the present study is built.

The practical difficulty is that the components involved belong to different administrative domains. Group membership is maintained by directory administrators, network policy is maintained by server administrators, and VLAN enforcement is maintained by network engineers. Vendor and platform documentation describes each element in isolation, but an integrated, measured account of the complete chain, from a directory group through a RADIUS policy to an authorised switch port, is much harder to obtain. That chain is what this study implements and evaluates.

1.1. Related work and the research gap

The normative basis for the architecture is well established. IEEE Std 802.1X-2020 specifies port-based network access control and the authenticator state machine [1]; RFC 3748 specifies EAP [2]; RFC 2865 specifies RADIUS and assigns UDP port 1812 for authentication [3]; and RFC 2868 specifies the tunnel attributes that carry a VLAN assignment [7]. RFC 2759 specifies the MS-CHAPv2 exchange used as the inner method in this study [8], and RFC 5216 specifies EAP-TLS, the certificate-based alternative [9]. RFC 5176 defines the dynamic authorisation extensions that allow an already-authorised session to be re-authorised or disconnected [10], a capability that is relevant to the extensions discussed later in this paper but that was not exercised here.

Platform documentation covers the server side of the deployment. Microsoft documents the NPS role and its policy model [5], recommended NPS configuration practice [6], the EAP methods available for network access in Windows [11], and the deployment of 802.1X authenticated access using NPS together with Active Directory Certificate Services [12]. On the network side, vendor documentation describes dynamic VLAN assignment driven by the RADIUS tunnel attributes; On the network side, vendor documentation provides worked examples of dynamic VLAN assignment using RADIUS tunnel attributes, including wireless LAN controller deployments [13]. The present study applies the same RADIUS attribute mechanism to a wired IEEE 802.1X authenticator, where the returned VLAN authorisation is enforced at the access port. Mares Mendoza reports a wireless deployment in which VLANs are assigned dynamically from directory attributes [14], which is the closest published counterpart to the work reported here. NIST Special Publication 800-207 situates this class of control within the broader movement away from implicit trust based on network location [15].

Two gaps follow from this literature. First, the sources are component-oriented: each describes one link of the chain correctly, but none reports the behaviour of the assembled chain, including what

happens when it fails. Second, the closest empirical counterparts are wireless, whereas the wired access layer is where port-based control was originally defined and where static port-to-VLAN binding remains the prevailing practice. The contribution of this study is therefore not a new protocol or algorithm, but an integrated wired implementation whose functional behaviour, including its behaviour under authentication failure and under authentication-server unavailability, is measured and reported with quantified uncertainty.

1.2. Research questions and contributions

The study addresses three questions.

1. Can Active Directory security-group membership be translated into network-level authorisation by NPS and RADIUS, and enforced as a VLAN assignment by a wired IEEE 802.1X authenticator?
2. What authentication latency does the assembled chain exhibit, and how precisely can that latency be characterised from the number of trials performed?
3. does the architecture contain failed authentication and authentication-server unavailability without granting access to the production VLANs?

The contributions are correspondingly threefold. The study reports a complete and reproducible configuration of the directory, policy and network layers required for identity-based VLAN assignment on a wired access port; it reports the functional outcome of every trial with exact binomial confidence intervals, so that the strength of the evidence is visible rather than implied; and it characterises the two containment paths, the Auth-Fail VLAN and critical authentication, which determine whether the design fails safely.

2. Methods

2.1. Research design

The study uses a single-testbed functional verification design. A complete NAC architecture was built in a virtualised laboratory, five scenarios were defined to exercise its authorisation, rejection and failure paths, and each scenario was executed five times consecutively. The primary outcome for every trial is categorical: whether the switch port was placed in the VLAN that policy required. Authentication latency was recorded as a secondary, descriptive outcome. The design is deliberately deterministic rather than comparative: no experimental factor is randomised, no condition is withheld, and no claim of statistical inference between conditions is made.

2.2. Testbed architecture

The laboratory was implemented on a VMware Workstation host with GNS3 providing network emulation. The Windows Server virtual machine carried the Active Directory Domain Services, Active Directory Certificate Services and Network Policy Server roles for the domain abdulraheem.com. A domain-joined Windows 10 virtual machine acted as the IEEE 802.1X supplicant. A Cisco IOSvL2 15.2 image acted as the authenticator and RADIUS client, and a Cisco 3725 image provided inter-VLAN routing. The GNS3 virtual machine used two adapters, with adapter 1 connected to VMnet8 for NAT/internet access and adapter 2 connected to VMnet1 as the host-only laboratory segment. The Windows Server and Windows 10 client virtual machines were connected to VMnet1 for laboratory traffic. Table 1 lists the components and their configured values, and Fig. 1 shows the resulting logical topology.

Table 1. Components and configured values of the virtualised testbed.

Element	Configured value or product	Function in the testbed
Hypervisor	VMware Workstation	Hosts the server and client virtual machines

Network emulation	GNS3 virtual machine	Interconnects the virtual machines with the emulated network devices
Server	Windows Server	Domain controller, certification authority and RADIUS server
Client	Windows 10	Domain-joined IEEE 802.1X supplicant
Directory service	Active Directory Domain Services	Holds user accounts and the security groups used for authorisation
Authentication server	Microsoft Network Policy Server	Evaluates connection request and network policies
Certification authority	Active Directory Certificate Services	Issues the server certificate used by PEAP
AAA transport	RADIUS (UDP 1812 and UDP 1813)	Carries authentication, authorisation and accounting
Port access control	IEEE 802.1X	Holds the access port unauthorised until authentication succeeds
Outer method	PEAP	Establishes the TLS tunnel that protects the inner exchange
Inner method	EAP-MSCHAPv2	Validates the user password inside the tunnel
Authenticator	Cisco IOSvL2 15.2	RADIUS client; applies the returned VLAN authorisation
Router	Cisco 3725	Provides inter-VLAN routing
Management VLAN	VLAN 99 (MGMT), 192.168.254.140/24	Switch management interface
Staff VLAN	VLAN 10 (STAFF)	Production VLAN authorised for members of VLAN10-Users
Student VLAN	VLAN 20 (STUDENTS)	Production VLAN authorised for members of VLAN20-Users
Guest VLAN	VLAN 30 (GUEST)	Configured on the switch but not exercised in the reported scenarios
Quarantine VLAN	VLAN 999 (QUARANTINE)	Containment VLAN for authentication failure and server unavailability

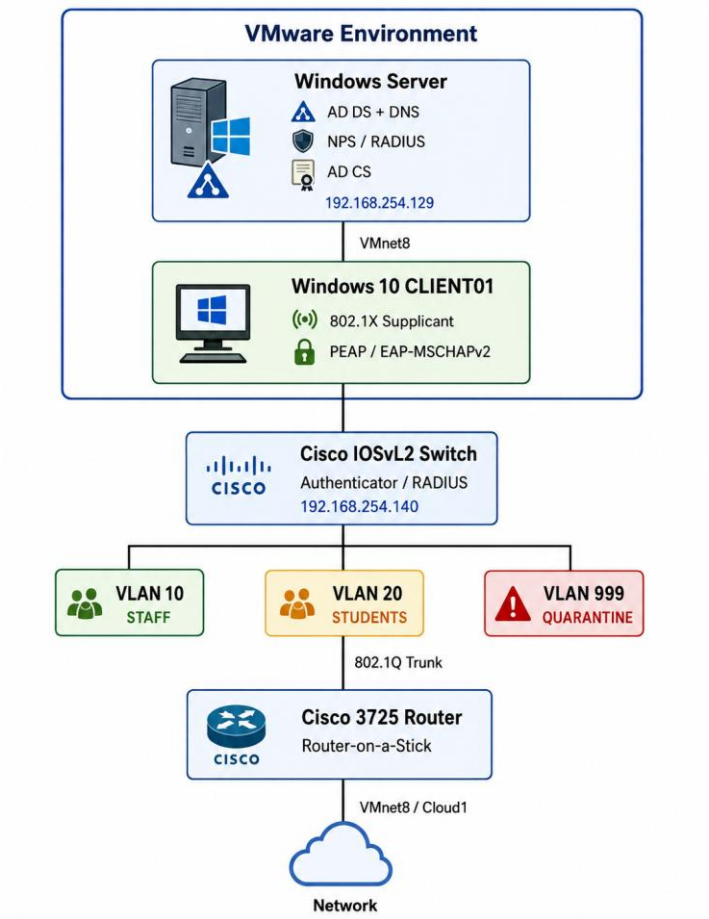


Fig. 1. Logical topology of the virtualised testbed, showing the supplicant, the IEEE 802.1X authenticator, the NPS/RADIUS server and the configured VLANs.

Alt text: Logical network topology showing a Windows 10 supplicant connected to a Cisco IOSvL2 authenticator, with Active Directory/NPS/RADIUS providing authentication and VLAN authorisation and a Cisco 3725 providing inter-VLAN routing.

2.3. Identity and authorisation design

Two security groups, VLAN10-Users and VLAN20-Users, were created in the domain, and one user account was placed in each. Group membership is the only user attribute that the authorisation decision depends upon, which keeps the mapping between identity and network position explicit and auditable. Table 2 records the mapping.

Table 2. Mapping from user account to security group and authorised VLAN.

User account	Active Directory security group	Authorised VLAN	VLAN name
ahmed	VLAN10-Users	10	STAFF
ali	VLAN20-Users	20	STUDENTS

NPS was installed and registered in Active Directory, and the switch was added as a RADIUS client using UDP 1812 for authentication and UDP 1813 for accounting. Two network policies were created. The first matches members of VLAN10-Users, grants access using PEAP and returns the tunnel attributes for VLAN 10; the second matches members of VLAN20-Users and returns the tunnel attributes for VLAN 20. The attribute triplet returned in each Access-Accept message is the one defined in RFC 2868 [7]:

```
Tunnel-Type                = VLAN     (13)
Tunnel-Medium-Type         = IEEE 802 (6)
Tunnel-Private-Group-ID    = 10   |   20
```

No policy returns a VLAN for any other user, so a request that matches no network policy is rejected rather than defaulted, and rejection is handled by the switch as described in Section 2.5. Fig. 2 shows the resulting authorisation chain from the supplicant through to the enforced VLAN.

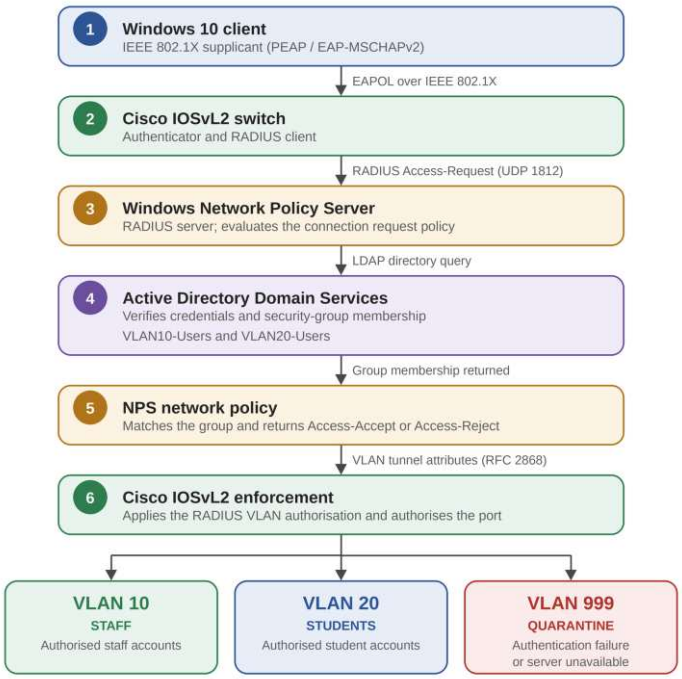


Fig. 2. The end-to-end authorisation chain. Group membership held in Active Directory is converted by NPS into RADIUS tunnel attributes and enforced by the authenticator as a VLAN assignment.

Alt text: Authorisation chain showing Active Directory group membership flowing to NPS policy evaluation, then through RADIUS tunnel attributes to Cisco switch VLAN enforcement.

2.4. Authentication method and certificate trust

Active Directory Certificate Services was installed as an enterprise root certification authority, and a server certificate was issued to the NPS server and bound to the PEAP configuration. The Windows 10 client was configured to enable IEEE 802.1X on the wired interface, to use PEAP as the outer method, to validate the NPS server certificate against the enterprise root, to use EAP-MSCHAPv2 [8] as the inner method, and to present the Windows logon credentials automatically. Server-certificate validation is what helps prevent an unauthenticated device from impersonating the authentication server and harvesting the inner exchange; Fig. 3 shows the trust path. PEAP with EAP-MSCHAPv2 is treated throughout this paper as the method under evaluation, not as a recommendation for high-assurance deployment, and the certificate-based alternative is discussed in Section 4.3.

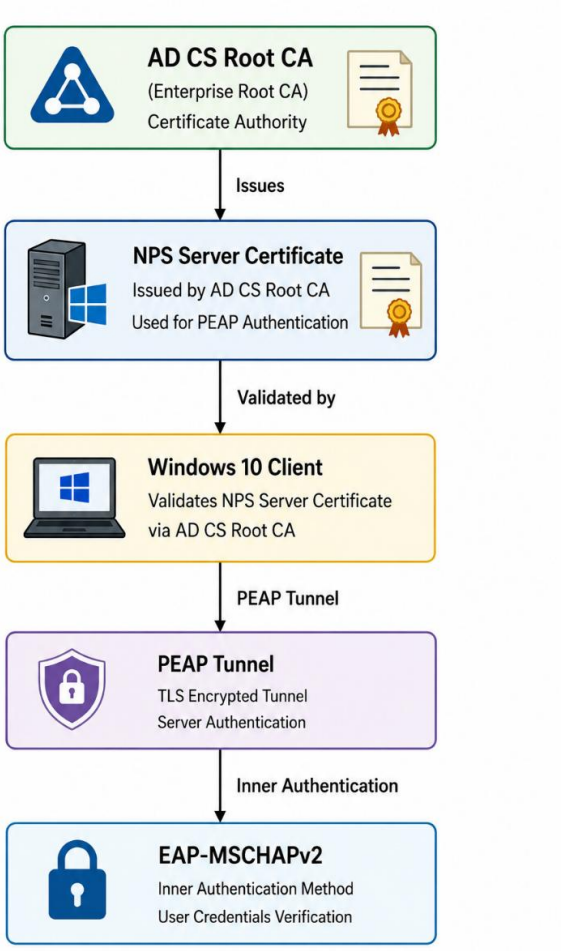


Fig. 3. Certificate trust path for PEAP. The client validates the NPS server certificate against the enterprise root certification authority before the inner EAP-MSCHAPv2 exchange begins.

Alt text: PEAP certificate trust path showing the Windows client validating the NPS server certificate against the enterprise root certification authority before EAP-MSCHAPv2 authentication.

2.5. Authenticator configuration

VLANs 10, 20, 30, 99 and 999 were created on the switch, VLAN 99 carrying the management interface and VLAN 999 serving as the containment VLAN. AAA was enabled with aaa new-model, with IEEE 802.1X authentication and authorisation directed to the RADIUS server group and accounting enabled. Client-facing interfaces were configured as access ports running the authenticator state machine. The two event actions in the listing below are the mechanisms that determine the failure behaviour reported in Sections 3.3 and 3.4: the first covers a rejected authentication and the second covers an unreachable authentication server, a behaviour also documented by the vendor as inaccessible authentication bypass.

```
! management interface
interface vlan 99
 ip address 192.168.254.140 255.255.255.0
 no shutdown
!
! access port facing the supplicant
authentication port-control auto
authentication host-mode single-host
authentication order dot1x
authentication priority dot1x
dot1x pae authenticator
authentication event fail action authorize vlan 999
authentication event server dead action authorize vlan 999
no cdp enable
```

2.6. Experimental scenarios

Five scenarios were defined to exercise the authorisation path, the rejection path and the server-failure path. Each was executed five times consecutively, giving 25 authentication events in total. Table 3 states the condition and the policy-required outcome for each scenario; the observed outcomes are reported in Section 3.

Table 3. Experimental scenarios, conditions and policy-required outcomes.

Scenario	Condition applied	Outcome required by policy	Trials
T1	ahmed with valid credentials	Port authorised into VLAN 10	5
T2	ali with valid credentials	Port authorised into VLAN 20	5
T3	Valid account with an incorrect password	Access rejected; port placed in VLAN 999	5
T4	Account matching no authorising network policy	Access rejected; port placed in VLAN 999	5
T5	NPS service stopped; RADIUS authentication service unavailable	Server marked dead; port placed in VLAN 999	5

2.7. Measurement procedure

For each trial the scenario identifier, the authentication outcome, the VLAN in which the port was placed and the authentication latency were recorded. Authentication latency was measured at the authenticator interface from the timestamp of the EAPOL-Start frame that begins the exchange to the timestamp at which the port reached the authorised state. The interval therefore contains every stage of the chain, not the RADIUS round trip alone: it includes EAPOL initiation, the EAP identity exchange, the RADIUS Access-Request, NPS policy processing, Active Directory validation, the RADIUS Access-Accept carrying the tunnel attributes, application of the VLAN authorisation and final port authorisation. Fig. 4 sets out those stages, and they should be borne in mind when the latency values in Section 3.2 are compared with figures reported elsewhere, because a measurement covering only the RADIUS exchange is not comparable with one covering the whole chain.

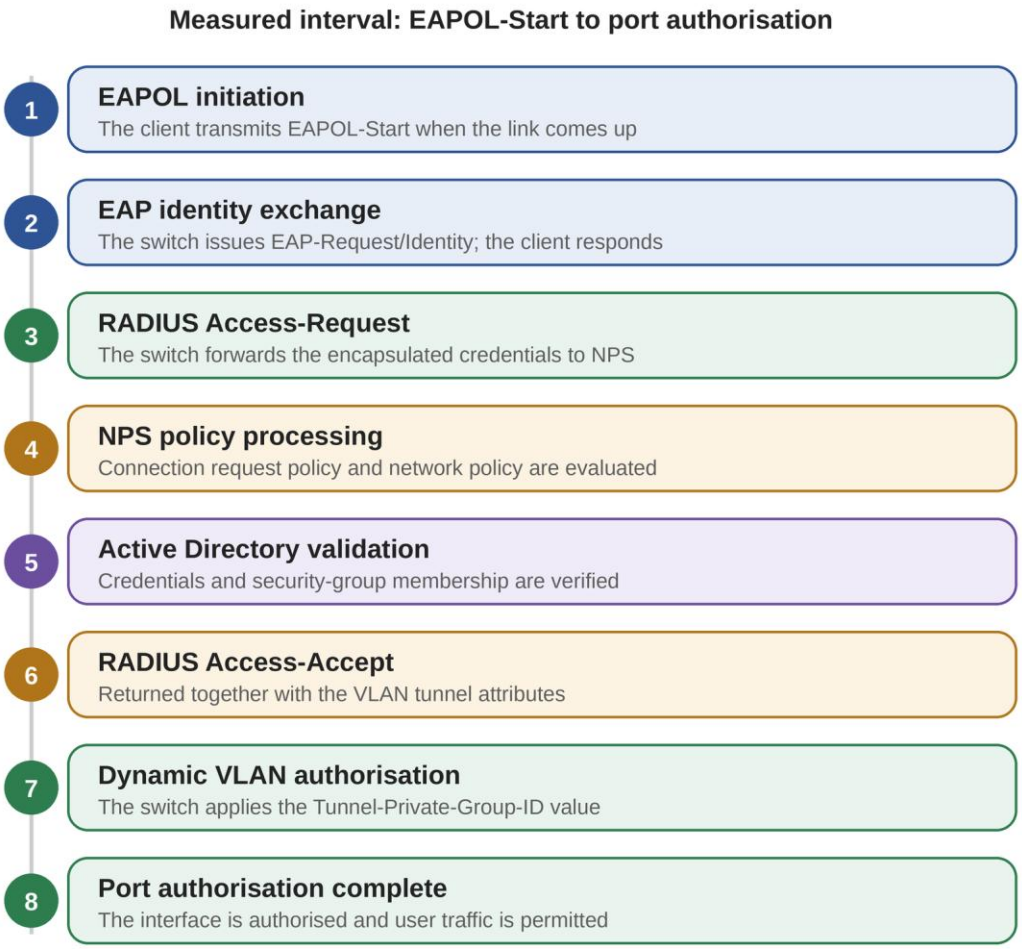


Fig. 4. Stages contained within the measured authentication latency, from EAPOL-Start to port authorisation.

Alt text: Authentication-latency sequence from EAPOL initiation through EAP identity exchange, RADIUS request, NPS processing, Active Directory validation, Access-Accept, VLAN application, and final port authorisation.

2.8. Sample size and statistical analysis

Five trials per scenario were performed. This number was chosen because the primary outcome is deterministic rather than stochastic: for a given configuration the authorisation decision is produced by a protocol state machine and a policy evaluation, so repeated trials test the stability and reproducibility of that decision rather than sampling a distribution of possible decisions. Repetition therefore guards against transient faults and configuration accidents, and it is reported here as such.

The limits of what five trials can establish are stated explicitly rather than left implicit. Every proportion in Section 3 is reported as a count out of the number of trials and is accompanied by an exact Clopper–Pearson 95% confidence interval, computed from the beta distribution, which is the interval appropriate to small binomial samples and does not rely on a normal approximation. Five successes out of five trials yield a point estimate of 100% with a 95% confidence interval of 47.8% to 100%; the interval, not the point estimate, is the honest summary of the evidence. Achieving an exact 95% lower confidence bound of 95% would require 72 consecutive successes, and a lower bound of 99% would require 368. Wilson score intervals were also computed and are reported alongside the exact intervals for comparison.

Latency is treated as descriptive throughout. Dispersion is reported as the observed minimum and maximum rather than as a standard deviation, because a standard deviation estimated from five observations is itself very imprecise: its 95% confidence interval extends from 0.60 to 2.87 times the sample value, so quoting it to two decimal places would imply a precision the design cannot support. No inferential comparison is made between the latencies of the two authorisation scenarios; Section 4.1 gives the analysis that supports that decision. Confidence intervals and the supporting computations were produced with a purpose-written routine implementing the incomplete beta function, the Clopper–Pearson and Wilson intervals, the chi-square interval for a standard deviation and Welch's t test. The statistical calculations reported in this section were checked against the recorded trial data in Supplementary Data 2. The complete measurement protocol and the completed per-trial dataset are provided as Supplementary Data 1 and 2, respectively.

3. Results

3.1. Authentication and authorisation outcomes

Across the 25 authentication events the port was placed in the VLAN required by policy in every case. Table 4 reports each scenario as a count with its exact and Wilson 95% confidence intervals. In the two authorisation scenarios the supplicant was authenticated and the port was authorised into the VLAN corresponding to the user's security group, confirming that group membership was carried end to end and applied by the authenticator. In the two rejection scenarios and in the server-failure scenario the port was placed in VLAN 999, and no trial in any scenario resulted in access to VLAN 10 or VLAN 20 by an endpoint that policy did not authorise for it.

Table 4. Observed outcomes with exact (Clopper–Pearson) and Wilson 95% confidence intervals. The intervals are wide because five trials per scenario were performed; they describe the precision of the evidence, not variability in the observed behaviour.

Scenario and measured quantity	Count	Rate	Exact 95% CI	Wilson 95% CI
T1 authentication success	5/5	100%	47.8–100%	56.6–100%
T1 assignment to VLAN 10	5/5	100%	47.8–100%	56.6–100%
T2 authentication success	5/5	100%	47.8–100%	56.6–100%
T2 assignment to VLAN 20	5/5	100%	47.8–100%	56.6–100%
T3 rejection of an incorrect password	5/5	100%	47.8–100%	56.6–100%
T4 rejection of an unauthorised account	5/5	100%	47.8–100%	56.6–100%
T3 and T4 pooled: containment in VLAN 999	10/10	100%	69.2–100%	72.2–100%
T5 containment in VLAN 999 when NPS is unreachable	5/5	100%	47.8–100%	56.6–100%
Production-VLAN exposure in T3, T4 and T5	0/15	0%	0–21.8%	0–20.4%
All scenarios pooled: outcome required by policy	25/25	100%	86.3–100%	86.7–100%

3.2. Authentication latency

Table 5 reports the latency observed in the two authorisation scenarios and the two rejection scenarios. Mean authentication latency was 1.42 s for VLAN 10 and 1.48 s for VLAN 20, with observed ranges of 1.25 s to 1.68 s and 1.28 s to 1.74 s respectively. The overall mean across the two authorisation scenarios was 1.45 s. Rejection and containment completed in a mean of 0.85 s in both T3 and T4; the observed

ranges were 0.80–0.90 s and 0.82–0.87 s respectively. Rejection and containment are shorter than successful authorisation because the exchange terminates at the Access-Reject message without the subsequent production-VLAN authorisation stages shown in Fig. 4.

Table 5. Observed authentication latency. Dispersion is reported as the observed range; see Section 2.8 for why no standard deviation is quoted at this sample size.

Scenario	Trials	Mean	Observed minimum	Observed maximum
T1 authorisation into VLAN 10	5	1.42 s	1.25 s	1.68 s
T2 authorisation into VLAN 20	5	1.48 s	1.28 s	1.74 s
T1 and T2 combined	10	1.45 s	1.25 s	1.74 s
T3 rejection and containment	5	0.85 s	0.80 s	0.90 s
T4 rejection and containment	5	0.85 s	0.82 s	0.87 s

The two authorisation means differ by 0.06 s, or 4.1% of the smaller value. The observed ranges overlap over almost their entire extent, from 1.28 s to 1.68 s. The two rejection scenarios both have a mean of 0.85 s, but their observed ranges differ slightly. No difference between the two authorisation scenarios is claimed, and Section 4.1 quantifies why the design cannot support such a claim.

3.3. Containment of failed authentication

In scenario T3 an existing account was presented with an incorrect password, and in scenario T4 an account was presented that matched no authorising network policy. In both scenarios NPS returned an Access-Reject message and the authenticator applied the configured Auth-Fail action, placing the port in VLAN 999. Containment occurred in 10 of 10 pooled trials, an exact 95% confidence interval of 69.2% to 100%. The distinction between the two scenarios is worth noting because they exercise different decision points: T3 fails inside Active Directory credential validation, whereas T4 fails at network-policy matching in NPS. Both nevertheless converge on the same enforcement action at the authenticator, which is the property that makes the containment behaviour predictable. Fig. 5 sets out the sequence.

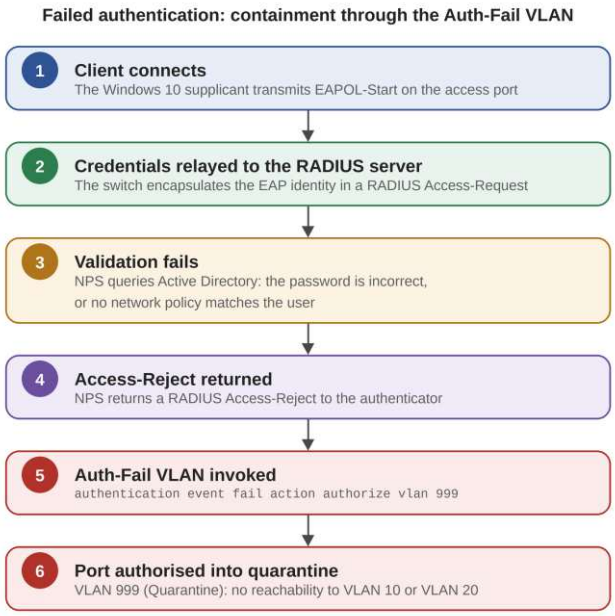


Fig. 5. Containment of a failed authentication. The Access-Reject message triggers the configured Auth-Fail action and the port is authorised into VLAN 999 rather than into a production VLAN.

Alt text: Failed-authentication containment sequence in which an Access-Reject causes the Cisco authenticator to apply VLAN 999 instead of a production VLAN.

3.4. Behaviour when the authentication server is unreachable

In scenario T5, the NPS service was stopped, making the RADIUS authentication service unavailable while the supplicant attempted to authenticate. The authenticator retransmitted the Access-Request, marked the RADIUS server as dead once its configured criteria were met, and applied the configured critical-authentication action, placing the port in VLAN 999. Table 6 reports the observed behaviour and Fig. 6 sets out the sequence. The total observed interval from the start of the exchange to port authorisation in VLAN 999 was 4.82 s under the configured timeout, retransmission, and dead-server criteria. This value is therefore a configuration- and implementation-dependent observation rather than a general performance characteristic of the architecture.

Table 6. Observed behaviour when the authentication server is unreachable.

Measured quantity	Observed value
Trials	5
RADIUS authentication service state	Unavailable (NPS service stopped)
Configured RADIUS timeout	3 s
Total interval to authorisation in VLAN 999	4.82 s
Fallback VLAN applied	999 (QUARANTINE)
Containment in VLAN 999	5/5; exact 95% CI 47.8–100%
Access to VLAN 10 or VLAN 20	0/5; exact 95% CI 0–52.2%

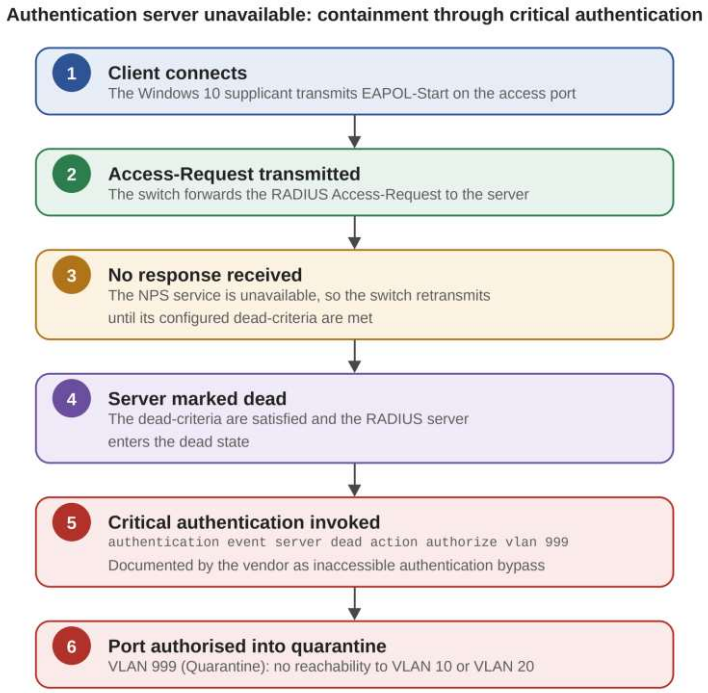


Fig. 6. Containment when the authentication server is unreachable. Once the dead-server criteria are satisfied, critical authentication places the port in VLAN 999.

Alt text: Authentication-server failure sequence in which the RADIUS server becomes unreachable, the authenticator marks it dead, and critical authentication places the endpoint in VLAN 999.

3.5. Dynamic compared with static VLAN assignment

Table 7 contrasts the implemented architecture with static port-to-VLAN binding across the properties that the testbed exercised. The comparison is qualitative and architectural; no static configuration was measured against the dynamic one in this study, and the table should be read accordingly.

Table 7. Architectural comparison of static VLAN binding with the implemented dynamic assignment.

Property	Static VLAN binding	Implemented dynamic assignment
Basis of the VLAN decision	Physical switch port	Authenticated identity and policy
Authentication before access	Not inherent	Required by IEEE 802.1X
Directory integration	None	Authorisation derived from security-group membership
RADIUS required for VLAN assignment	No	Yes
Effect of moving a user to another port	VLAN changes with the port	VLAN follows the user
Changing a user's authorisation	Reconfiguration of the switch port	Group or policy change on the server
Point of administration	Each access switch	Directory and policy server
Handling of a failed authentication	Typically none	Automatic containment in VLAN 999
Behaviour if the policy server is unreachable	Unaffected	Containment in VLAN 999 by critical authentication

4. Discussion

4.1. Interpretation of the findings

The first research question is answered affirmatively and without qualification as to mechanism: in every trial, a security group held in Active Directory determined the VLAN into which a wired access port was placed. The chain functioned as designed, and the two failure paths behaved as configured. The appropriate qualification concerns strength of evidence rather than mechanism. Twenty-five consecutive correct outcomes support an exact 95% confidence interval of 86.3% to 100% for the probability of a correct authorisation decision; the design demonstrates that the mechanism works and is reproducible, and it does not establish a reliability figure of the kind an operator would need before deployment at scale.

The second question concerns latency. The measured interval of approximately 1.4 s to 1.5 s spans the entire chain of Fig. 4, and it is small enough that an interactive user would not perceive port

authorisation as a delay. The 0.06 s difference between the two authorisation scenarios should not be interpreted as a difference between them. The two authorisation scenarios differed in mean latency by only 0.06 s (1.42 s versus 1.48 s). Because the study was designed for functional verification rather than statistical comparison of latency between conditions, and because only five trials were performed per scenario, this difference is treated as descriptive rather than inferential. The overlapping observed ranges further indicate that no comparative latency conclusion is warranted from the present dataset.

The third question concerns fail-safe behaviour, and it is the result of greatest practical interest. Access control that fails open is worse than no access control, because it produces a false expectation of segmentation. Two distinct failure modes were exercised. A rejected authentication is a decision the policy server actively makes and communicates, whereas an unreachable policy server means no decision is available at all, and the authenticator must act on a locally configured default. Both were contained in VLAN 999 in every trial. The second case is the more consequential, because the fallback behaviour of an authenticator whose authentication server has become unreachable is entirely a matter of local configuration: the same event that produced containment here would, with a different configured action, have produced unrestricted access on the port. The result is therefore properly read as evidence that the mechanism is available and effective when configured, and as an argument for treating that configuration as a required control rather than an optional refinement.

4.2. Relation to previous work

The behaviour observed here is consistent with the normative specifications: the tunnel attributes of RFC 2868 were honoured by the authenticator [7], the EAP exchange was transported within RADIUS as specified in RFC 3579 [4], and the authenticator state transitions followed IEEE Std 802.1X-2020 [1]. The study's contribution relative to the platform documentation [5], [6], [11], [12] is that it reports the assembled chain rather than its components, including the two failure paths, which the component documentation describes individually but does not evaluate together. Relative to the closest empirical counterparts, which are wireless deployments [13], [14], the present study confirms that the same attribute-driven authorisation operates on a wired access port, where static port-to-VLAN binding remains common practice and where the physical-access assumption that motivates port-based control is at its weakest. The architecture also illustrates in miniature the principle set out in NIST Special Publication 800-207, that network position should follow verified identity rather than confer implicit trust [15], although the present design applies that principle only at the moment of port authorisation and not continuously through the life of a session.

4.3. Security considerations

Three properties of the implemented design deserve explicit comment. First, authorisation is centralised: a change of a user's network position is made by editing a group membership or a network policy on the server, and requires no change to any switch configuration, which removes a class of configuration drift in which stale port assignments outlive the people they were created for. Second, containment is automatic rather than administrative: a failed authentication produces a defined network position instead of an ambiguous one, which both limits exposure and preserves the connectivity needed for remediation. Third, and less favourably, the evaluated authentication method is password-based. EAP-MSCHAPv2 inside a PEAP tunnel depends for its security on the client validating the server certificate, because an inner exchange conducted with an unvalidated server is exposed; the client in this testbed was configured to perform that validation, but the requirement is a configuration discipline rather than a protocol guarantee. EAP-TLS [9] removes the shared secret from the exchange altogether by authenticating both parties with certificates, at the cost of a certificate lifecycle for every endpoint, and it is the appropriate direction for a deployment with higher assurance requirements.

The architecture also leaves two gaps that follow from its scope rather than from any defect in configuration. Authorisation is evaluated once, at port authorisation; an endpoint whose posture or

entitlement changes afterwards retains its VLAN until the session ends, unless the dynamic authorisation extensions of RFC 5176 [10] are deployed to re-authorise or disconnect sessions in flight. And single-host mode, used here, ties the authorisation decision to the first supplicant on the port, so a device placed behind an authenticated endpoint is not independently authenticated.

4.4. Limitations

The limitations of this study are substantial and are stated plainly, because they bound what the results can be used for.

- Sample size. Five trials per scenario support the conclusion that the mechanism functions reproducibly, and they do not support a reliability estimate. The exact confidence intervals in Table 4 are the correct summary of the evidence: five successes out of five bound the true success probability only at 47.8% to 100%.
- Virtualised environment. VMware and GNS3 introduce host scheduling, emulated forwarding and virtual switching into every measured interval. The latency values in Table 5 characterise this testbed and are not transferable to physical hardware.
- Scale. The testbed contained one authenticator, one authentication server and one supplicant. Contention at the policy server, the behaviour of multiple simultaneous authentications and the effect of a redundant authentication server were not exercised.
- Latency dispersion. Per-trial latency values are provided in Supplementary Data 2, but the sample size is small. The paper therefore reports the mean and observed range rather than making distributional or inferential claims from five observations per scenario. No standard deviation is reported.
- Configuration dependence of the failure timing. The 4.82 s interval in Table 6 is determined by the configured RADIUS timeout and dead-criteria and would change with those settings; it is not a property of the architecture.
- Absence of instrument-level evidence. The results are reported as measurements taken at the authenticator interface. Packet captures, authenticator session output and NPS event records were collected during the trials according to the evidence protocol in Supplementary Data 1, but the individual capture and log files are not included in this submission. The reader cannot independently verify individual trials from the material presented here.
- Single method under evaluation. Only PEAP with EAP-MSCHAPv2 was tested. EAP-TLS, machine authentication, posture assessment and guest onboarding were configured or considered but not evaluated.

4.5. Future work

The most immediate extension is a larger and instrumented trial programme: enough repetitions to give the success-rate intervals practical width, automated capture and archival of per-trial latency and packet-level evidence, and publication of archived packet captures as supplementary data so that individual trials can be audited. Beyond that, four directions follow directly from the limitations above. Replacing PEAP with EAP-TLS [9] would remove the password from the exchange and allow the two methods to be compared on both security and latency. Deploying a redundant authentication server would allow the failure scenario to be re-run against a design that fails over rather than falls back, which is the behaviour an operator would actually want. Introducing the dynamic authorisation extensions of RFC 5176 [10] would convert the one-time authorisation decision into a revisable one, and combining that with endpoint posture assessment would allow authorisation to reflect device state as well as user identity. Finally, extending the testbed to multiple authenticators and a realistic endpoint population would allow the scaling behaviour of the policy server, and the effect of simultaneous authentications, to be measured rather than assumed.

5. Conclusions

This study implemented and evaluated a complete identity-based network access control architecture on a wired access layer, integrating Active Directory Domain Services, Active Directory Certificate Services, Microsoft Network Policy Server, RADIUS, IEEE 802.1X with PEAP and EAP-MSCHAPv2, and Cisco IOSvL2 switching, in a virtualised VMware and GNS3 laboratory. Across 25 authentication events in five scenarios, the switch port was placed in the VLAN required by policy in every case: members of VLAN10-Users were authorised into VLAN 10 and members of VLAN20-Users into VLAN 20, failed authentications were contained in VLAN 999, and the same containment applied when the authentication server was made unreachable. No trial produced access to a production VLAN for an endpoint that policy did not authorise for it. Pooled over all scenarios, the proportion of outcomes required by policy was 25 of 25, an exact 95% confidence interval of 86.3% to 100%.

Mean authentication latency, measured from EAPOL-Start to port authorisation, was 1.42 s for VLAN 10 (observed range 1.25 s to 1.68 s) and 1.48 s for VLAN 20 (observed range 1.28 s to 1.74 s), with rejection and containment completing in a mean of 0.85 s. The 0.06 s difference between the two authorisation scenarios is not statistically supportable at this sample size and no difference is claimed.

The substantive conclusion is that directory group membership can be carried, without loss, through a RADIUS policy decision into an enforced VLAN assignment on a wired port, and that the implemented configuration contained both evaluated failure modes by placing the endpoint in VLAN 999 following a rejected authentication and following authentication-server unavailability. The evidence is drawn from a single small-scale virtualised testbed and is reported with confidence intervals that make its precision explicit; it establishes feasibility and reproducibility rather than operational reliability. On that basis the testbed is offered as a reproducible model for teaching, for network access control experimentation and as a starting point for the larger-scale evaluation that a deployment decision would require.

6. Ethics statement

Not applicable. This study used a virtualised laboratory testbed and did not involve human participants, animals, human-derived research materials, or personally identifiable research data.

7. Conflicts of interest

The author declares no conflicts of interest relevant to this work.

8. CRediT authorship contribution statement

Abdulraheem Hamoud Mohammed Almaweri: Conceptualization; methodology; software; investigation; data curation; formal analysis; writing—original draft; writing—review and editing; visualisation; project administration.

9. AI-assisted tools disclosure

AI-assisted tools were used for language editing and manuscript preparation. The author reviewed and verified all technical content, experimental records, analyses, citations, and the final manuscript.

10. Funding

This work was supported by no external funding. It was conducted as an independent academic research project in a virtualised laboratory environment.

11. Acknowledgements

The author thanks the Academy of Industrial Science and Technology, Djibouti, for providing an academic environment that supports research and practical experimentation in computer networking and cybersecurity.

12. Supplementary material

The measurement protocol and reproduction procedure (Supplementary Data 1), together with the completed per-trial dataset (Supplementary Data 2), are included in this manuscript. Supplementary Data 2 contains the 25 valid trial records underlying the results reported in this paper.

13. Supplementary Material S1 — Measurement Protocol

Measurement protocol and reproduction procedure for "Experimental Implementation of Network Access Control with Dynamic VLAN Assignment Using Active Directory, Windows NPS, RADIUS, and Cisco 802.1X"

This document specifies how the testbed was assembled, how each trial was run, how the authentication interval was timed, what evidence a complete trial record contains, what invalidates a trial, and how the reported results were computed from the trial records. It accompanies Supplementary Data 2, the completed per-trial workbook containing the 25 valid trial records used to produce the results reported in the paper.

13.1. Scope and definitions

Table S1.1. Terms used throughout the paper and this protocol.

Term	Definition used throughout the paper and this protocol
Trial	One authentication attempt by one supplicant on one access port, from a port in the unauthorised state through to a recorded outcome.
Scenario	A defined condition applied to a trial (T1 to T5 in Table 3 of the paper). Each scenario was executed five times consecutively.
Policy-required outcome	The VLAN in which the port must end up if the configured policy is working, decided before the trial is run, not after.
t_0	The instant the authentication exchange starts on the access port. See Section 6.
t_1	t_1 — The instant the authenticator reaches the final enforcement state for the trial: either successful authorisation into the policy-required production VLAN or configured containment in VLAN 999. See Section 6.
Authentication interval	$t_1 - t_0$. This is the quantity reported as "authentication latency" in Section 3.2 of the paper. It spans the whole chain of Fig. 4, not the RADIUS round trip alone.
Containment	The port being placed in VLAN 999 (Quarantine) with no reachability to VLAN 10 or VLAN 20.
Void trial	A trial that failed the validity rules of Section 8. It is recorded as void with its reason and repeated; it is never silently discarded.

13.2. Software and images

Record the exact version or build of every item.

Table S1.2. Software and image versions to be recorded.

Role	Product used	Version or build to record

Hypervisor	VMware Workstation	17.5.1
Network emulation	GNS3	GNS3 2.2.54 + GNS3 VM 2.2.54
Server operating system	Windows Server	Windows Server 2019 — Build 17763
Client operating system	Windows 10	Windows 10 22H2 — Build 19045
Authenticator image	Cisco IOSvL2	Version 15.2(20200924:215240)
Router image	Cisco 3725	IOS Version 12.4(15)T14

13.3. Build sequence

13.3.1. Host and network emulation

1. Install the hypervisor and GNS3, and import the IOSvL2 and 3725 images.
2. Give the GNS3 virtual machine two adapters: adapter 1 connected to VMnet8 (NAT) for internet access, and adapter 2 connected to VMnet1 (host-only) for laboratory traffic.
3. Attach the server and client virtual machines to the host-only laboratory segment (VMnet1) and confirm they can reach the emulated switch.
4. Record the resources assigned to each virtual machine (vCPU, RAM).

Virtual Machine	vCPU	RAM
Windows Server	2	4 GB
Windows 10 Client	2	4 GB
GNS3 VM	4	8 GB

13.3.2. Directory services

5. Promote the server to a domain controller for the domain `abdulraheem.com` and record the functional level.
6. Create the two security groups `VLAN10-Users` and `VLAN20-Users`.
7. Create one user account in each: `ahmed` in `VLAN10-Users`, `ali` in `VLAN20-Users`. Group membership is the only user attribute the authorisation decision depends on, so keep the accounts otherwise identical.
8. For scenario T4, a third user account was created that was not a member of either VLAN10-Users or VLAN20-Users. The account name used was testuser.
9. Join the Windows 10 client to the domain and verify the domain logon of both accounts before any 802.1X configuration is applied.

13.3.3. Certificate services and PEAP trust

1. Install Active Directory Certificate Services as an enterprise root certification authority.
2. Issue a server certificate to the NPS host from the “RAS and IAS Servers” template and confirm that its enhanced key usage includes Server Authentication and that its subject matches the name the client will validate.
3. Confirm the enterprise root certificate is present in the client’s Trusted Root Certification Authorities store — by domain membership it should be, because a client configured to validate a certificate it cannot chain will fail in a way that looks like a credential problem.
4. Record the certificate thumbprint, validity dates and template used.

13.3.4. Network Policy Server

1. Install the Network Policy and Access Services role and register the server in Active Directory.
2. Add the switch as a RADIUS client using its management address 192.168.254.140, with a shared secret. Record that a shared secret was configured; do not record its value.

3. Confirm the authentication and accounting ports are UDP 1812 and UDP 1813 respectively, and that they match the switch configuration.
4. Create the connection request policy that handles the requests locally.
5. Create two network policies. The first grants access to members of `VLAN10-Users`; the second to members of `VLAN20-Users`. Set the EAP type to PEAP with EAP-MSCHAPv2 as the inner method in both.
6. On each policy return exactly the RFC 2868 triplet, with the group identifier differing between the two:

```
Tunnel-Type               = VLAN(13)
Tunnel-Medium-Type        = IEEE 802(6)
Tunnel-Private-Group-ID   = 10|20
```

7. Verify the policy processing order, and verify that no policy grants access to any other user. A request that matches no policy must be rejected rather than defaulted — that is what scenario T4 tests.
8. Enable Network Policy Server accounting so that every decision is logged with a timestamp.

13.3.5. Authenticator

The configuration below is the one reported in Section 2.5 of the paper, with the RADIUS and timing settings the protocol requires added explicitly. Substitute your own addresses and interface identifiers.

```
       ! --- logging with millisecond timestamps: needed for the timing method
service timestamps log datetime msec localtime
service timestamps debug datetime msec localtime
logging buffered 64000 debugging
ntp server 192.168.254.129
!
! --- VLANs
vlan 10
 name STAFF
vlan 20
 name STUDENTS
vlan 30
 name GUEST
vlan 99
 name MGMT
vlan 999
 name QUARANTINE
!
interface vlan 99
 ip address 192.168.254.140 255.255.255.0
 no shutdown
!
! --- AAA and RADIUS
aaa new-model
radius server NPS
 address ipv4 192.168.254.129 auth-port 1812 acct-port 1813
 key <shared-secret>
!
aaa group server radius RAD-NPS
 server name NPS

       !
aaa authentication dot1x default group RAD-NPS
aaa authorization network default group RAD-NPS
aaa accounting dot1x default start-stop group RAD-NPS
dot1x system-auth-control
!
```

```
       ! --- retransmission and dead-server criteria: whole seconds only.
!     Record these values; the T5 interval is a function of them.
radius-server timeout 3
radius-server retransmit 3
radius-server dead-criteria time 3 tries 3
radius-server deadtime 5
!
! --- access port facing the supplicant
interface FastEthernet0/1
 switchport mode access
 authentication port-control auto
 authentication host-mode single-host
 authentication order dot1x
 authentication priority dot1x
 dot1x pae authenticator
 authentication event fail action authorize vlan 999
 authentication event server dead action authorize vlan 999
 no cdp enable
```

13.3.6. Supplicant

1. Set the Wired AutoConfig service to start automatically and start it.
2. On the Ethernet adapter, enable IEEE 802.1X authentication and select PEAP as the network authentication method.
3. In the PEAP properties, enable verification of the server certificate, select the enterprise root certification authority, and set the inner authentication method to EAP-MSCHAPv2.
4. Enable automatic use of the Windows logon name and password so that the credential presented is the domain credential and no interactive prompt enters the measured interval.
5. Record the adapter model as presented to the guest and any driver-level power-saving settings, which can delay link-up and therefore delay to.

13.4. Pre-trial verification

Complete all checks below before beginning the trials.

Table S1.3. Pre-trial checks, all of which must pass before a trial counts.

Check	How	Expected
VLANs exist	show vlan brief	10, 20, 30, 99 and 999 all present and active
802.1X globally enabled	show dot1x all	System auth control enabled; the access port listed as an authenticator
RADIUS reachable	show aaa servers	The NPS entry present, state up, no dead counters
End-to-end AAA	test aaa group RAD-NPS ahmed <password> new-code	Authentication successful
NPS decision logged	Windows Security log on the server	An event for the test above, with the matched connection request policy and network policy named
Certificate trust	Client certificate stores	The enterprise root present under Trusted Root Certification Authorities
Clocks aligned	show clock detail and the server clock	Synchronise the switch clock with the NPS/Windows Server time source and record the observed offset in seconds.
Port starts unauthorised	show authentication sessions interface Fa0/1 details	No session, or status unauthorised, before each trial begins

13.5. Per-trial procedure

Steps common to every trial

1. Confirm the port is in the unauthorised state and clear any residual session: clear authentication sessions interface FastEthernet0/1.
2. Clear the log buffer so the trial's messages stand alone: clear logging.
3. Start the capture on the link between the supplicant and the access port, before the link comes up.
4. Apply the scenario condition from Table S1.4 below.
5. Bring the link up — administratively with shutdown followed by no shutdown, or by reconnecting the adapter in the guest. Use the same method for every trial.
6. Wait until the port reaches a stable state, then stop the capture.
7. Collect the evidence set of Section 7 and enter the trial in Supplementary Data 2 while the evidence is in front of you.
8. Restore the testbed to its baseline — in particular restart the NPS service after every T5 trial and confirm the RADIUS server is marked alive again before the next trial starts.

Table S1.4. Condition applied and restoration required for each scenario.

Scenario	Condition applied before the link comes up	Policy-required outcome	Restoration afterwards
T1	Log on as ahmed, valid password	Port authorised into VLAN 10	None
T2	Log on as ali, valid password	Port authorised into VLAN 20	None
T3	Valid account, deliberately incorrect password	Access rejected; port placed in VLAN 999	Confirm the account is not locked out; record the lockout policy
T4	Account belonging to neither authorising group	Access rejected; port placed in VLAN 999	None
T5	Stop the Network Policy Server service so the RADIUS server is unreachable	Server marked dead; port placed in VLAN 999	Restart the service; confirm the server is alive again

13.6. Timing method

Two independent timings are taken for every trial in T1, T2, T3 and T4, and both are recorded.

For T3 and T4, t_1 denotes the final containment state in VLAN 999 rather than successful production-VLAN authorisation.

They should agree; if they do not, the discrepancy is itself a finding and belongs in the record.

13.6.1. Primary timing — the authenticator's own clock

Both endpoints of the interval are taken from the switch's log, so no cross-host clock skew enters the measurement at all:

- **t_0** — the millisecond-stamped entry recording that the authentication method started on that interface, of the form `%AUTHMGR-5-START: Starting 'dot1x' for client … on Interface Fa0/1`.
- **t_1** — the entry recording the final enforcement outcome for that client, either successful production-VLAN authorisation or the configured containment action in VLAN 999. Because message strings may differ between IOS releases, record the verbatim log line emitted by the image and identify the message used as the t_1 endpoint.

13.6.2. Cross-check — the link capture

From the capture taken on the access link, take t_0 as the timestamp of the EAPOL-Start frame and t_1 as the timestamp of the first frame the port forwards on the assigned VLAN after the RADIUS Access-Accept. Record the capture-derived interval alongside the log-derived one. Where the two differ by more than one logging interval, report the capture-derived value and note the difference in the trial record.

13.6.3. What the interval contains, and what it does not

The interval spans the complete authentication and authorisation sequence represented in Fig. 4 of the paper: EAPOL initiation, the EAP identity exchange, the RADIUS Access-Request, NPS policy processing, Active Directory validation, the Access-Accept carrying the tunnel attributes, application of the VLAN authorisation, and final port authorisation. It is therefore not comparable with a figure measured across the RADIUS exchange alone, and any comparison with published latencies must establish that the same span was measured.

The virtualised environment may introduce host-scheduling effects; therefore, latency values are reported as characteristics of the experimental testbed rather than physical hardware.

13.6.4. Scenario T5 is a configuration measurement, not a protocol one

In T5 no server responds, so the interval is determined by radius-server timeout, radius-server retransmit and radius-server dead-criteria, plus the time the authenticator takes to apply the critical-authentication action. Record all four configured values with the measurement and report them together, as Table 6 of the main manuscript does. Changing any of them changes the result; the number describes a configuration, not the architecture.

13.7. Evidence set for one trial

A trial record is complete only when every item below has been stored under the trial identifier. Adopt one naming convention and keep it: `T<scenario>-<trial>-<yyyymmdd-hhmmss>`, for example `T1-03-20260826-141207.pcapng`.

Table S1.5. Evidence set retained for a single trial.

Item	Where it comes from	What it establishes
Port state and assigned VLAN	show authentication sessions interface Fa0/1 details	The primary outcome: which VLAN the port was actually placed in
VLAN membership of the port	show vlan brief	Corroborates the assignment independently of the session table
RADIUS server state	show aaa servers	Whether the server was alive or dead during the trial — essential for T5
Log excerpt containing t_0 and t_1	show logging	The primary timing, and the authenticator's own account of the decision
NPS decision record	Windows Security log on the server	The event identifier, the matched connection request policy and network policy, and the attributes returned
Link capture	Wireshark on the access link	The cross-check timing and the EAPOL and RADIUS exchange itself
Reachability test result	A ping or equivalent from the client to a host in VLAN 10 and in VLAN 20	That containment in VLAN 999 was real and not merely labelled

Client-side state	The adapter's authentication status on the client	Corroborates the outcome from the endpoint's point of view

13.8. Trial validity rules

A trial is void, and must be repeated, if any of the following holds:

- the port had not returned to the unauthorised state before the trial began;
- the capture was not running for the whole interval;
- the switch log is missing either t_0 or t_1;
- the server's Security log contains no matching decision within the trial window;
- any configuration on the switch, the server or the client was changed during the trial;
- the host was under a load that visibly stalled the virtual machines;
- the RADIUS server was still marked dead from a previous T5 trial.

Void trials are entered in the workbook with the void flag set and the reason recorded, and the count of void trials is kept. Discarding trials without recording them is what turns a small sample into a biased one, and it is not recoverable after the fact.

13.9. Analysis procedure

1. Enter every valid trial as one row of the Trials sheet of Supplementary Data 2.
2. The Summary sheet computes, per scenario, the count of policy-required outcomes k out of n, the exact Clopper–Pearson 95% interval from the beta distribution, and the Wilson score interval.
3. For latency, report the mean with the observed minimum and maximum and state n. Do not report a standard deviation at n = 5: its own 95% interval runs from 0.60 to 2.87 times the sample value, so two decimal places would imply a precision the design cannot support.
4. No inferential comparison was performed between the latency of T1 and T2 because the study was designed for functional verification rather than comparative performance testing, and only five trials were conducted per scenario.
5. Before any numerical result enters the paper, reproduce the calculation independently from the recorded S2 trial data and confirm that the result agrees with the reported value. Every interval quoted in Tables 4 and 6 of the paper was checked against the recorded trial data.
6. Check the arithmetic of every summary against the algebraic bound $\max|x_i - \bar{x}| / s \le (n - 1)/\sqrt{n}$. A summary that violates it cannot have come from real observations, and a reviewer who checks will find it. The check was applied to the recorded trial data.

13.10. Reporting conventions

- Every rate appears as a count out of the number of trials, with an exact confidence interval beside it.
- Dispersion appears as the observed minimum and maximum with n stated, never as a standard deviation from a handful of observations.
- Any timing that depends on configured timers is labelled as such and is reported together with the configured values.
- No difference is claimed that the design cannot support, however suggestive the point estimates look.
- What was observed but not archived is stated plainly rather than implied to be available.

13.11. Scaling the protocol

Section 4.5 of the paper recommends a larger instrumented trial programme. Table S1.6 states what "larger" has to mean to buy a given claim, so that the next round is sized before it is run rather than after.

Table S1.6. Trials required to support a given claim.

Claim to be supported	Requirement

Exact 95% lower bound of 80% on the success rate	17 consecutive successes
Exact 95% lower bound of 90%	36 consecutive successes
Exact 95% lower bound of 95%	72 consecutive successes
Exact 95% lower bound of 99%	368 consecutive successes

The first four rows follow from the exact binomial interval and assume an unbroken run of successes: a single failure raises the requirement substantially. All six figures were independently checked from the reported sample-size and statistical assumptions.

14. Supplementary Material S2 — Per-Trial Dataset

Trial ID	Scen.	Date	Time started	User name	Req. VLAN	Obs. VLAN	Auth result	Interval used (s)	Void?	Notes
T1-01	T1	01/03/26	14:10	ahmed	10	10	Accept	1.25	No	Baseline trial for VLAN 10
T1-02	T1	01/03/26	14:15	ahmed	10	10	Accept	1.37	No	Normal authenticatio n chain
T1-03	T1	01/03/26	14:20	ahmed	10	10	Accept	1.4	No	Normal authenticatio n chain
T1-04	T1	01/03/26	14:25	ahmed	10	10	Accept	1.4	No	Normal authenticatio n chain
T1-05	T1	01/03/26	14:30	ahmed	10	10	Accept	1.68	No	Upper limit observed range
T2-01	T2	01/03/26	14:40	ali	20	20	Accept	1.28	No	Baseline trial for VLAN 20
T2-02	T2	01/03/26	14:45	ali	20	20	Accept	1.43	No	Normal authenticatio n chain
T2-03	T2	01/03/26	14:50	ali	20	20	Accept	1.47	No	Normal authenticatio n chain
T2-04	T2	01/03/26	14:55	ali	20	20	Accept	1.49	No	Normal authenticatio n chain
T2-05	T2	01/03/26	15:00	ali	20	20	Accept	1.74	No	Upper limit observed range
T3-01	T3	01/03/26	15:10	ahmed	999	999	Reject	0.8	No	Incorrect password attempt
T3-02	T3	01/03/26	15:15	ahmed	999	999	Reject	0.83	No	Access-Reject returned by

										NPS
T3-03	T3	01/03/26	15:20	ahmed	999	999	Reject	0.85	No	Access-Reject returned by NPS
T3-04	T3	01/03/26	15:25	ahmed	999	999	Reject	0.87	No	Access-Reject returned by NPS
T3-05	T3	01/03/26	15:30	ahmed	999	999	Reject	0.9	No	Access-Reject returned by NPS
T4-01	T4	01/03/26	15:40	testuser	999	999	Reject	0.82	No	No network policy matched
T4-02	T4	01/03/26	15:45	testuser	999	999	Reject	0.84	No	No network policy matched
T4-03	T4	01/03/26	15:50	testuser	999	999	Reject	0.85	No	No network policy matched
T4-04	T4	01/03/26	15:55	testuser	999	999	Reject	0.87	No	No network policy matched
T4-05	T4	01/03/26	16:00	testuser	999	999	Reject	0.87	No	No network policy matched
T5-01	T5	01/03/26	16:10	ahmed	999	999	No response	4.82	No	NPS Service stopped, fallback
T5-02	T5	01/03/26	16:15	ahmed	999	999	No response	4.82	No	NPS Service stopped, fallback
T5-03	T5	01/03/26	16:20	ahmed	999	999	No response	4.82	No	NPS Service stopped, fallback
T5-04	T5	01/03/26	16:25	ahmed	999	999	No response	4.82	No	NPS Service stopped, fallback
T5-05	T5	01/03/26	16:30	ahmed	999	999	No response	4.82	No	NPS Service stopped, fallback

15. References

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

[2] B. Aboba, L. Blunk, J. Vollbrecht, J. Carlson and H. Levkowetz, Eds., “Extensible Authentication Protocol (EAP),” Internet Engineering Task Force, RFC 3748, Jun. 2004.

[Online]. Available: https://www.rfc-editor.org/rfc/rfc3748.html. [Accessed: Aug. 26, 2026].

[3] C. Rigney, S. Willens, A. Rubens and W. Simpson, “Remote Authentication Dial In User Service (RADIUS),” Internet Engineering Task Force, RFC 2865, Jun. 2000. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2865.html. [Accessed: Aug. 26, 2026].

[4] B. Aboba and P. Calhoun, “RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP),” Internet Engineering Task Force, RFC 3579, Sep. 2003. [Online]. Available: https://www.rfc-editor.org/rfc/rfc3579.html. [Accessed: Aug. 26, 2026].

[5] Microsoft, “Network Policy Server (NPS) overview,” Microsoft Learn. [Online]. Available: https://learn.microsoft.com/en-us/windows-server/networking/technologies/nps/nps-top. [Accessed: Aug. 26, 2026].

[6] Microsoft, “Network Policy Server best practices,” Microsoft Learn. [Online]. Available: https://learn.microsoft.com/en-us/windows-server/networking/technologies/nps/nps-best-practices. [Accessed: Aug. 26, 2026].

[7] G. Zorn, D. Leifer, A. Rubens, J. Shriver, M. Holdrege and I. Goyret, “RADIUS Attributes for Tunnel Protocol Support,” Internet Engineering Task Force, RFC 2868, Jun. 2000. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2868.html. [Accessed: Aug. 26, 2026].

[8] G. Zorn, “Microsoft PPP CHAP Extensions, Version 2,” Internet Engineering Task Force, RFC 2759, Jan. 2000. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2759.html. [Accessed: Aug. 26, 2026].

[9] D. Simon, B. Aboba and R. Hurst, “The EAP-TLS Authentication Protocol,” Internet Engineering Task Force, RFC 5216, Mar. 2008. [Online]. Available: https://www.rfc-editor.org/rfc/rfc5216.html. [Accessed: Aug. 26, 2026].

[10] M. Chiba, G. Dommety, M. Eklund, D. Mitton and B. Aboba, “Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS),” Internet Engineering Task Force, RFC 5176, Jan. 2008. [Online]. Available: https://www.rfc-editor.org/rfc/rfc5176.html. [Accessed: Aug. 26, 2026].

[11] Microsoft, “Extensible Authentication Protocol (EAP) for network access,” Microsoft Learn. [Online]. Available: https://learn.microsoft.com/en-us/windows-server/networking/technologies/extensible-authentication-protocol/network-access. [Accessed: Aug. 26, 2026].

[12] Microsoft, “Deploy 802.1X authenticated wired or wireless access,” Microsoft Learn Core Network Guide. [Online]. Available: https://learn.microsoft.com/en-us/windows-server/networking/core-network-guide/cncg/wireless/a-deploy-8021x-wireless-access. [Accessed: Aug. 26, 2026].

[13] Cisco Systems, “Configure a RADIUS server and WLC for dynamic VLAN assignment,” Cisco Technical Documentation. [Online]. Available: https://www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-vlan/71683-dynamicvlan-config.html. [Accessed: Aug. 26, 2026].

[14] E. Mares Mendoza, “Red inalámbrica con asignación de VLAN dinámica,” Cuadernos Técnicos Universitarios de la DGTIC, vol. 2, no. 3, 2024, doi: 10.22201/dgtic.ctud.2024.2.3.66.

[15] S. Rose, O. Borchert, S. Mitchell and S. Connelly, “Zero Trust Architecture,” National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST Special Publication 800-207, Aug. 2020, doi: 10.6028/NIST.SP.800-207.