

Article

SNMP Exposure Within Corporate Networks

Pelayo Nuño , David Álvarez , Alicia Flórez , Francisco G. Bulnes  and Juan C. Granda 

Department of Computer Science, University of Oviedo, Campus de Viesques, 33204 Gijón, Asturias, Spain; uo237670@uniovi.es (D.Á.); florezalicia@uniovi.es (A.F.); bulnes@uniovi.es (F.G.B.); jcgranda@uniovi.es (J.C.G.)

* Correspondence: nunopelayo@uniovi.es

Abstract

The Simple Network Management Protocol (SNMP) remains the de facto standard for monitoring corporate networks. However, many devices lack support for its latest version, and common administrator misconfigurations expose organizations to severe security risks. This study aims to provide guidance on improving SNMP configuration security by identifying devices exposed through legacy SNMPv1 and SNMPv2c within private corporate networks. A grey-box penetration testing methodology was conducted across 78 public and private organisations in Spain, employing the scanning tools *Onesixtyone* and *Metasploit* to identify active devices utilising well-known community strings. The assessment reveals widespread vulnerabilities: more than 90% of the analysed organisations have assets exposed via well-known read-only community strings, and over 60% are vulnerable through well-known read-write community strings, resulting in a total of 2403 and 645 vulnerable devices, respectively. This exposure is highly critical in network hardware, affecting over 50% (read-only) and nearly 40% (read-write) of the organisations. Furthermore, multiple active disclosed vulnerabilities using *OpenVAS* and prevalent administrative flaws were identified. These findings underscore that basic deployment misconfigurations leave extensive internal assets exposed, highlighting an urgent need for structural remediation and improved administrative training within corporate environments.

Keywords: SNMP; networks; exposure; security; management

1. Introduction

The Simple Network Management Protocol (SNMP) is an application-layer protocol within the OSI model, widely used for managing Internet Protocol (IP) networks [1]. SNMP employs a centralised architecture consisting of an SNMP manager and one or more managed devices, referred to as SNMP agents. The software agent is embedded in the managed device to allow remote administration, enabling the SNMP manager to control and monitor the network. The manager can read values from an agent using GET operations or modify them through SET operations. The manager relies on definitions in the Management Information Base (MIB) to perform operations on the managed device, such as retrieving variable values and processing events (TRAP operations).

Although the SNMP architecture faces scalability and real-time adaptability challenges in dynamic, decentralized networks [2], it remains the standard for monitoring corporate networks [3]. However, despite its ubiquity in corporate environments, SNMP can introduce severe threats to critical infrastructures [4]. This vulnerability stems from its susceptibility to brute-force, Denial of Service (DoS), Distributed Denial of Service (DDoS), and reflection attacks, among others [5].



Academic Editor: Hakim Mellah

Received: 17 July 2026

Revised: 10 September 2026

Accepted: 14 September 2026

Published: 21 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

SNMP has evolved through three versions: SNMPv1, SNMPv2, and SNMPv3. Although SNMPv3 provides enhanced security and authentication [1], earlier versions remain widely supported for compatibility and preferred for their simplicity [6]. SNMPv1 and SNMPv2c use the so-called community strings to read (read-only) or modify (read-write) information on a managed device. These community strings are included in every SNMP message and transmitted in plaintext. Moreover, many network device manufacturers configure default community strings, such as the well-known read-only *public* and read-write *private* variants, which numerous administrators continue to use in production environments [6]. Even when personnel attempt to configure new community strings, these custom credentials are often insufficiently secure and remain vulnerable to simple guessing attacks. This situation is even more critical when network administrators are non-specialists, lacking proper training in both network management and security, and are unaware of the presence of SNMP on their devices or its default configuration and behaviour.

Nowadays, security incidents occur with high frequency and often result in the exposure of private and sensitive data. The root causes of such incidents are rarely sophisticated attacks. Instead, they stem from simple misconfigurations on devices [7]. Studies have shown that the most frequent personal factor contributing to these misconfigurations by administrators is a lack of knowledge [8]. Consequently, insufficient administrator training often delays the detection of security misconfigurations, a problem that is continuously amplified by an expanding attack surface [9]. Furthermore, such vulnerabilities are often prevalent and severe, and are typically detected during testing or, in the worst case, during deployment and operation [10]. Indeed, empirical evidence indicates that organisations frequently operate with persistent network misconfigurations across their infrastructure [11]. In the case of SNMP, this implies that numerous devices exposed across public and private networks have SNMP enabled and use well-known community strings, without the responsible administrator even being aware of it.

The main objective of this work is to provide guidance to organisations on improving their SNMP configuration security. This study analyses SNMPv1 and SNMPv2c exposure within private corporate networks. Specifically, the study focused on 78 organisations in Spain, encompassing networks of varying sizes and spanning both public and private sectors. In each organisation, a grey-box penetration test was conducted, simulating an attack to identify SNMP-managed devices with well-known community strings. For this purpose, the scanning tool *Onesixtyone* was employed [12]. Although several studies have examined SNMP exposure in public IP networks [13–16], to the best of the authors' knowledge, this is the first study to address such exposure within the internal private networks of organisations.

The assessment reveals widespread vulnerabilities, demonstrating that a vast majority of the analysed organisations have assets exposed via well-known read-only credentials, while a substantial proportion remains vulnerable through read-write strings. This exposure is particularly critical within core network hardware. Beyond empirical quantification, the analysis identifies multiple active SNMP-related vulnerabilities alongside prevalent administrative misconfigurations, highlighting systemic weaknesses in the protocol's deployment across these environments.

The remainder of the paper is organised as follows. Section 2 provides a brief overview of SNMP, including its architecture, functionality, and security concerns. Related work on SNMP exposure and misconfigurations is reviewed in Section 3. The research methodology, including the analysis of the participating organisations, is presented in Section 4, while the results obtained from the study are detailed in Section 5. Section 6 offers a discussion of the findings and their implications. Finally, Section 7 concludes the paper and outlines directions for future work.

2. Background

SNMP is a general-purpose network management protocol that enables remote management of any type of device. A typical SNMP-based network management system follows a client–server architecture, where a manager manages many devices. For a device to be managed remotely using SNMP, it must have management software installed and running, commonly referred to as an SNMP agent.

An SNMP agent is internally composed of three elements, as shown in Figure 1. The management interface defines the types of messages that an agent processes and the events it generates. The agent logic processes incoming requests, translating them into actions on the managed device as requested by the manager, while the MIB serves as the database containing the information needed to manage the device. Depending on the nature of the device, the information contained in the MIB varies. Thus, the MIB of a printer may store information such as the remaining ink or paper, while the MIB of a network switch stores information about port status.




Figure 1. Architecture of an SNMP agent.

The information contained in an MIB is organised hierarchically, and each piece of management information is assigned an Object Identifier (OID). Each time a manager sends a management request (such as changing the device name or disabling an interface), it specifies the OID of the piece of information it is attempting to read or modify.

This simple architecture relies on clearly defined MIBs and a limited set of message types. As a result, SNMP agents are extremely lightweight and can run even on resource-constrained devices. This versatility has driven their widespread adoption across the industry.

Security Concerns in SNMP

SNMP exhibits several fundamental flaws inherent to its protocol design, as it lacks adequate mechanisms for mutual authentication between agents and managers and operates over the unreliable User Datagram Protocol (UDP) [17]. It also suffers from well-known vulnerabilities in certain flawed SNMP implementations [18], which can be exploited through integer and buffer overflow attacks, among others [19].

SNMPv1 and SNMPv2c are vulnerable to sniffing attacks because they use unencrypted community strings for authentication. Attackers can use automated tools to obtain these strings. Dictionary attacks can also test well-known community strings, as agents are often configured with default values, such as *public* (read-only) and *private* (read-write). If the read-only community string is set to *public*, an attacker can retrieve device information

via SNMP, including hostnames, interface and routing information, uptime, Transmission Control Protocol (TCP) session details, open ports, and other network attributes [4]. A more severe threat arises when the read–write community string is set to *private*. This configuration enables direct modification of device settings, allowing attackers to change configurations and gain full control over the device.

As noted above, SNMP uses UDP at transport level, which exposes it to vulnerabilities enabling DoS and DDoS attacks. A common vector is UDP flooding [20], which involves sending a large number of UDP packets from spoofed source addresses. When the manager receives such traffic, it attempts to process each request. Due to the high volume and UDP's lack of overload protection, the manager can be overwhelmed, preventing it from processing legitimate SNMP traffic [17].

In a related scenario, an attacker can send an SNMP message from a spoofed source address using a guessed read-only community string. The agent then responds to the third-party with the spoofed IP address, a behavior known as an SNMP reflection attack. In SNMPv2c, the *GetBulkRequest* enables the efficient retrieval of multiple MIB objects. A maliciously crafted *GetBulkRequest* may elicit responses containing entire tables, resulting in a substantial volume of data. An attacker can exploit this behaviour by sending spoofed *GetBulkRequest*'s, causing oversized responses to be reflected towards a victim, enabling DoS or DDoS attacks [21]. In fact, it has been shown that SNMP is one of the most commonly used protocols for conducting reflection attacks [22].

SNMPv3 was introduced to enhance the security of its predecessors, particularly focusing on authentication. However, SNMPv3 allows fingerprinting by malicious actors [23]. A manager must interact with an SNMPv3 agent to discover its unique engine ID before it can authenticate the session. SNMPv3 agents respond to unauthenticated engine ID discovery requests (REPORT PDUs) revealing this value. By default, the disclosed engine ID encodes device-specific information such as the Media Access Control (MAC) address, vendor identifier, or model details. This information leaks valuable system attributes, allowing attackers to conduct effective reconnaissance. In fact, the study in [24] shows that as few as ten SNMPv3 probe packets per router IP are sufficient to accurately fingerprint nearly two-thirds of routers in the IPv4 Internet.

SNMPv3 use in wireless networks has also posed security risks [25]. It relies on an administrator password for one-way authentication, from which encryption and authentication keys are derived. If an attacker obtains this password, they can potentially control all devices within the management domain associated with the compromised credentials. Furthermore, the one-way authentication mechanism can enable man-in-the-middle (MiTM) attacks.

Finally, any SNMP version can be exploited to perform a Cross-Channel Scripting (XCS) attack, a variant of Cross-Site Scripting (XSS) where a web application vulnerability allows the injection of malicious code via network protocols. The authors in [26] demonstrate XCS attacks through SNMP on embedded servers in devices such as cameras, wireless routers, and access points. These devices provide web interfaces enabling administrators to perform management tasks through a web browser. If SNMP agents are misconfigured with well-known read–write community strings and without Access Control Lists (ACLs), an attacker can use SNMP SET commands to write values to OIDs. If these values are later displayed on the web interface without proper escaping, they may contain malicious code that triggers an XCS attack when an administrator accesses the web page.

To mitigate such protocol-level weaknesses, recent architectural and technical proposals have aimed to enhance SNMP security and message integrity [27]. Nevertheless, as observed in this study, the practical security of SNMP deployments remains heavily compromised by fundamental administrative misconfigurations.

Based on the above, SNMP should be disabled entirely if not required. When its use is necessary, SNMPv3 should be preferred and treated as sensitive management traffic. Furthermore, the UDP ports assigned to SNMP (161 and 162) should be restricted to management networks only and must never be exposed to the Internet unless strictly necessary.

3. Related Work

Research on SNMP exposure has attracted significant attention in recent years. In this regard, numerous studies have examined the degree of SNMP exposure in public IP networks resulting from misconfigurations.

In [13], the authors analysed several protocols to assess security differences between IPv4 and IPv6. They examined SNMP exposure in dual-stack hosts identified via Domain Name System (DNS). They evaluated a total of 25K dual-stack routers and 520K dual-stack servers. For SNMP, two MIB objects (*sysDescr* and *sysName*) were retrieved via SNMPv2c requests using the read-only community string. Results showed that approximately 3% of routers and slightly more than 1% of servers had this community string configured with default values. SNMPv2c was also found to be over twice as exposed on routers over IPv6 compared to IPv4. Similar studies conducted during the same period reported more than 35K SNMP-exposed devices in the United States [28] and over 40K in Europe [29], including more than 10K detected in Germany [30]. The findings of these studies indicate that nearly all SNMP-exposed devices were routers, mainly operated by telecommunications companies. More recently, a study focused on the Nordic countries was conducted to determine how many SNMP-misconfigured devices could be used to perform reflection attacks [14]. Results showed approximately 13K devices discovered in Sweden, 4K in Norway, 3K in Denmark, and 1.5K in Finland. SNMP requests also confirmed that these reflectors correspond to home appliances, including networking equipment and smart Internet of Things (IoT) devices.

Building on the above, a relevant research area within public networks and the so-called observable Internet focuses on IoT networks. In this regard, vulnerabilities in Supervisory Control and Data Acquisition (SCADA) devices were assessed in [15]. Using *Nessus* as a scanning tool, more than 500K devices were discovered, with roughly 1.5K exhibiting high-severity vulnerabilities related to default community strings. These vulnerabilities affected devices from many major Programmable Logic Controller (PLC) manufacturers. In [31], the authors conducted a large-scale scan on IoT devices via *Shodan*, focusing on two consumer IoT types: smart TVs and printers. Results revealed 18.6K consumer IoT devices with default community strings. Specifically, 11.4K (61.6%) used the default read-only community string (*public*), while 7.5K (38.4%) used the default read-write community string (*private*).

Recent studies continue to confirm the widespread presence of misconfigured SNMP devices. In [23], a large-scale Internet-wide scan using SNMPv3 was performed. Over 4.6 million devices were fingerprinted, including approximately 350K routers. Specifically, the authors obtained detailed information regarding router vendor market share, deployment strategies of network operators worldwide, router uptime statistics, and the geographic distribution of vendors across different regions. Another study evaluating SNMP exposure using *Censys* datasets is presented in [32]. Results showed over 22 million routers discovered, with approximately 90% configured with SNMPv3, 5% with SNMPv2c, and the remaining 5% with SNMPv1. In addition, using *Shodan*, over 1K routers were found vulnerable to Common Vulnerabilities and Exposures (CVEs) involving DoS, data extraction, configuration manipulation, and remote code execution. Finally, the analysis of SNMP exposure was extended to publicly accessible devices affected by misconfigured firewalls in [16]. *Shodan* scans showed over 47K devices accessible via SNMPv1 or SNMPv2c

using the default read-only community string, including servers, routers, switches, printers, Voice over Internet Protocol (VoIP) devices, and others. Additionally, 465K devices could be fingerprinted using SNMPv3. Similarly, an alternative approach relying on *Nmap* to assess exposed targets behind firewalls is explored in [33].

The aforementioned studies share a common focus on the public IP domain, i.e., devices accessible over the Internet. In contrast, this work focuses on the internal corporate environment, analysing SNMP exposure within private networks, which could be exploited by intruders or attackers located inside the network. To the best of the authors' knowledge, this is the first study specifically addressing SNMP exposure in private corporate networks.

4. Methodology

This paper investigates the exposure of SNMPv1 and SNMPv2c in private corporate networks of Spanish organisations, covering entities of various sizes ranging from small to large. This section describes the classification criteria used for the organisational sample and provides a detailed explanation of the methodology employed in the study.

4.1. Organisations

A total of 78 Spanish organisations participated in this study. To enhance the representativeness of the results, the sample was categorised as shown in Table 1.

Table 1. Distribution of participating organisations by category.

Organisation Type	# of Organisations	Description
Local/provincial councils	27	Small, medium and large city councils together with provincial governments
Private enterprises	21	Organisations from diverse sectors, including retail, chemical manufacturing, the automotive industry, software development and information technology
Educational institutions	13	Universities and both public and private secondary schools
Public institutions	17	Agencies involved in healthcare, innovation, natural resource management, financial auditing and industrial regulation

Figure 2 shows the percentage distribution of organisations across the four categories, emphasising that public actors (councils, public institutions and some educational institutions) represent the main category of this study. Nevertheless, a substantial number of private actors (private enterprises and some educational institutions) were included in the analysis.

The organisations involved also differ considerably in scale and were grouped into four size-based categories according to the number of employees: 20 fall within the 0–100 employee range, 29 within 101–500, 18 within 501–2000, and the remaining 11 exceed 2000 employees. The distribution across these categories is illustrated in Figure 3, offering a representative cross-section of organisational sizes within the dataset. For ease of reference throughout the remainder of the paper, four corresponding labels (S1, S2, S3 and S4) are assigned to denote each size category.

A purposive sampling approach was employed to achieve a balanced representation across different categories and organisation sizes (based on employee count). Organisations were contacted directly to explain the study's objectives. Given the non-random nature of this sampling strategy, the results are presented in the subsequent section using descriptive statistics, as inferential statistical methods would not be appropriate for making population-

wide claims. All organisations provided informed consent and findings are reported anonymously to ensure confidentiality.




Figure 2. Percentage of organisations of each category.




Figure 3. Organisations by number of employees.

The research was conducted over five consecutive years, from 2022 to 2026. Figure 4 shows the annual number of organisations analysed. The highest concentration of data was collected during 2025 and the first half of 2026 (accounting for nearly half of the total sample), ensuring that the results remain current at the time of writing. Data were combined across the full five-year period to build a larger dataset. Although network configurations and vulnerabilities can change over time, SNMP misconfigurations may persist for years in real-world environments. Therefore, these results provide an overall overview across the study period rather than a single point in time or a temporal trend. In addition, data collection across organisation categories occurred continuously throughout the 2022–2026 period based on audit availability, ensuring that no single category was restricted to a specific timeframe.




Figure 4. Number of organisations in each year of research.

Furthermore, the sample exhibits geographical diversity, with organisations distributed across Spain to minimise territorial bias, as illustrated in Figure 5. Some Spanish regions, highlighted in red, have been excluded from the study, as the researchers were unable to establish contact with any willing participants in those areas. Nonetheless, the remaining regions are represented by at least one participating organisation, and most of them by several.




Figure 5. Distribution of organisations across the regions of Spain.

4.2. Method

The procedure followed for each organisation was the same. Organisations consented to grey-box penetration testing involving controlled malicious interactions with the network. In these analyses, SNMP exposure was identified through simulated attacks.

The methodology consisted of using the *Onesixtyone* scanning tool (version 0.7), available as part of the Kali Linux distribution, against a private network address segment to probe all devices within each subnet. The tool sends asynchronous SNMP queries (targeting the *sysDescr* OID) using a dictionary file (*Dict.txt*) (<https://github.com/trailofbits/>

[onesixtyone/blob/master/dict.txt](#) (accessed on 17 July 2026))) that contains common community strings. An example of the execution command using a generic private network segment and the target dictionary file is presented below:

```
root@kali:~# onesixtyone -c Dict.txt 192.168.1.0/24
```

A representative output resulting from a successful scan identifying active SNMP services and retrieving their respective *sysDescr* payload is as follows, where the identified community string is enclosed in square brackets:

```
192.168.1.13 [public] Vendor-A Router OS Version 15.0(2)
192.168.1.100 [1234] Vendor-B Managed Switch Revision 2.4.1
```

It is important to clarify that identifying a matching community string through dictionary testing encompasses two distinct operational scenarios: (i) factory-default configurations, where default community strings (e.g., *public* or *private*) remain active out of the box, and (ii) weak or common configurations, where administrators intentionally assign predictable, widely known strings. Throughout the rest of the article, the term ‘well-known strings’ is used to refer collectively to both scenarios. Distinctions between factory settings and manual administrator assignments are made explicitly when supported by vendor documentation or empirical evidence.

Once the initial scan concluded, each community string was tested on every device. Devices for which valid community strings were identified were subsequently subjected to further verification using the *Metasploit Framework* (version 6.1) modules *auxiliary/scanner/snmp/snmp_enum* and *auxiliary/scanner/snmp/snmp_set*. The purpose was to distinguish between read-only and read-write access permissions. Based on the previous hypothetical scan results, an example of read-only enumeration against host 192.168.1.13 using the identified *public* community string is presented below:

```
msf6 > use auxiliary/scanner/snmp/snmp_enum
msf6 auxiliary(scanner/snmp/snmp_enum) > set RHOSTS 192.168.1.13
msf6 auxiliary(scanner/snmp/snmp_enum) > set COMMUNITY public
msf6 auxiliary(scanner/snmp/snmp_enum) > run
```

Similarly, an example of write-permission verification against host 192.168.1.100 using the discovered ‘1234’ community string to update the *sysDescr* OID is shown as follows:

```
msf6 > use auxiliary/scanner/snmp/snmp_set
msf6 auxiliary(scanner/snmp/snmp_set) > set RHOSTS 192.168.1.100
msf6 auxiliary(scanner/snmp/snmp_set) > set COMMUNITY 1234
msf6 auxiliary(scanner/snmp/snmp_set) > set OID 1.3.6.1.2.1.1.1.0
msf6 auxiliary(scanner/snmp/snmp_set) > set OIDVALUE "New Value"
msf6 auxiliary(scanner/snmp/snmp_set) > run
```

To prevent operational disruption or persistent configuration changes on live production devices, write operations were conducted under a strict safety procedure. Specifically, the original OID value was read and logged prior to executing the *snmp_set* module. Verification was achieved by temporarily writing a new value. Immediately following successful confirmation of write access, a second *snmp_set* operation was executed to restore the original OID value to its initial state, thereby ensuring zero residual impact on the audited infrastructure.

Furthermore, to determine the presence of active, well-known CVE vulnerabilities, the *OpenVAS* vulnerability scanner (versions 21.4.4 to 23.50.0) was deployed, targeting the discovered SNMP-exposed assets of each organisation. Although several versions of the

tool were used, the involved plugins remained the same. Therefore, version changes had no impact on the study.

Finally, it should be noted that the entire networking environment of each organisation was not analysed. Access was limited to only part of the network, since internal routing between private subnets was sometimes restricted and only one or a few private network segments could be accessed physically. Data collection specifically targeted reachable operational and end-user subnets (such as student computer labs and common-area networks in educational institutions, general employee workstation segments in private enterprises, and public-access endpoints or administrative staff subnets in public institutions and local/provincial councils). This selection establishes an internal threat model that reflects a realistic operational attack surface. Specifically, it represents an adversary with a local network position, such as a non-privileged internal user, a compromised workstation, a rogue device, or a local visitor with physical access to an open port. While this partial coverage prevents definitive claims about the absolute security posture of an organisation, it renders our results in a conservative lower bound of total exposure. Nonetheless, the security of an organisation depends on its weakest component. Therefore, if any analysed segment revealed exposure, the organisation was classified as exposed, as it is vulnerable to the SNMP community-string scanning procedures described in this study.

5. Results


To provide a comprehensive overview, the results are organised into four subsections. First, exposure levels by organisation category are analysed in Section 5.1. Next, Section 5.2 examines the annual exposure prevalence across the five-year evaluation period. A quantitative and vendor breakdown of the identified assets is then provided in Section 5.3. Finally, the CVEs identified during the organisational network assessment, along with several prevalent misconfigurations and technical flaws, are analysed in Section 5.4.

5.1. Asset Exposure Analysis


Figure 6 presents the results of the analysis conducted across the studied organisations, according to their category. Specifically, Figure 6a shows the organisations where at least one device with a well-known read-only community string was discovered, whereas Figure 6b shows those where devices with a well-known read-write community string were present. As can be observed, the overall average exposure is high, since more than 91% of the analysed organisations (71) host at least one device responding to an SNMP GET operation via well-known read-only credentials.

Furthermore, 64% of the organisations (50) have devices vulnerable to unauthorised SNMP SET operations due to well-known read-write configurations. When examining specific categories, educational institutions exhibit the highest exposure rate, reaching 100% exposure via read-only community strings, closely followed by local and provincial councils. Conversely, private enterprises and public institutions present a lower proportion of positive cases, although their exposure rates remain remarkably high, exceeding 50% in both categories for read-write operations.

Next, the exposure level of network devices is evaluated. For the purpose of this specific analysis, the term *network devices* exclusively encompasses core networking and communication infrastructure components. This includes routers, Layer 2 and Layer 3 switches, Fibre Channel (FC) switches, firewalls, load balancers, Access Points (APs), and Wireless LAN controllers. Additionally, VoIP phones are also integrated into these criteria, as they are considered active network elements that effectively extend the operational perimeter of the access layer within a standard campus network architecture.



(a) Exposure via well-known SNMP read-only community strings




(b) Exposure via well-known SNMP read-write community strings


Figure 6. Percentage of organisations with exposed assets across the four analysed categories.

Figure 7 shows the results focusing solely on the aforementioned network devices, classified according to their organisational category. Specifically, Figure 7a shows the organisations where at least one network device was discovered with a well-known read-only community string, whilst Figure 7b shows those hosting network devices with well-known read-write communities. As observed in the overall average, 56% of the studied organisations (44) manage network infrastructure exposed via well-known read-only communities. More critically, 37% of these entities (29) fail to restrict well-known read-write communities on their networking core. When examining individual categories, private enterprises present the highest vulnerability rates, with 76% and 52% exposure in read-only and read-write permissions, respectively. Local and provincial councils follow closely as the second most impacted sector. Conversely, educational institutions exhibit a notably lower frequency rate in this specific layer, marking a distinct contrast to the global asset exposure trends observed previously.

These results highlight a notable contrast: private enterprises exhibit lower exposure rates globally, especially for read-write well-known communities, yet they exhibit the highest exposure level regarding core network devices. This finding is particularly concerning because SNMP exposure in routing, switching, and firewall infrastructure poses a significantly higher threat than in regular user devices. Compromising these core components grants attackers structural control over network segmentation and internal traffic routing, leaving corporate networks vulnerable. Conversely, the opposite is observed for educational institutions. While they emerged as the most vulnerable category during the global asset analysis, they exhibit the lowest exposure rates when focusing strictly on network devices.



(a) Exposure via well-known SNMP read-only community strings




(b) Exposure via well-known SNMP read-write community strings

Figure 7. Percentage of organisations with exposed network devices across the four analysed categories.

Finally, Figure 8 examines SNMP exposure according to organisation size. Overall, exposure levels increase alongside organisation size. The presence of at least one exposed asset via well-known read-only credentials (Column 1) is widespread, ranging from 80% in small organisations (S1) to 100% in the largest organisations (S4). Similarly, exposure via read-write credentials (Column 2) and exposures restricted to network devices (Columns 3 and 4) show higher prevalence in medium and large organisations (S2–S4) compared to small ones (S1).

Furthermore, category S3 (501–2000 employees) exhibits the highest rates for read-write exposures, reaching 89% for general assets (Column 2) and 67% for network devices (Column 4). In contrast, read-write exposure on network devices (Column 4) drops to 27% in S4 (>2000 employees) and 20% in S1 (0–100 employees), marking the lowest exposure values across all evaluated scenarios.

**Figure 8.** SNMP asset exposure prevalence by organisation size.

5.2. Annual Exposure Prevalence

Figure 9 provides a detailed breakdown of the annual prevalence of exposure to well-known SNMP community strings within the analysed organisations across the five-year data collection period.




Figure 9. SNMP asset exposure prevalence by year analyzed (2022–2026).

As can be observed, the exposure levels remain elevated throughout the entire observation period. Specifically, read-only exposure ranges between 82% (2026) and 95% (2025). Similarly, read-write exposure varies between 53% (2026) and 79% (2023). Although a lower percentage is observed in the final sampling period (2026), the multi-year empirical data indicates that well-known community strings remain present across the evaluated years, representing a systemic operational issue regarding SNMP management.

5.3. Quantitative and Vendor Breakdown

Table 2 outlines the breakdown of exposed devices utilising well-known read-only SNMP community strings, while Table 3 details those exposed via read-write strings. In total, 2403 devices were identified through well-known read-only strings (Table 2) and 645 through read-write strings (Table 3), resulting in a ratio of 3.73 read-only exposures for every read-write exposure. Note that this ratio reflects the overall prevalence of each exposure type (read-only vs. read-write) across the scanned infrastructure, rather than the simultaneous use of both community strings on a single device.

Because internal device inventories were not accessible, total baseline counts per device category across the target organisations were not recorded. Therefore, the relative frequencies in Tables 2 and 3 represent the proportion of each device type within the detected exposed sample ($N = 2403$ and $N = 645$, respectively), rather than comparative susceptibility rates relative to total deployed infrastructure. As shown in Table 2, network printers are the most exposed asset category for read-only operations, followed by print servers. Hosts and Layer 2 and Layer 3 switches also show notable exposure. Similarly, Table 3 highlights that network printers remain the primary asset category vulnerable to read-write operations, followed in this case by IoT sensors, Layer 2 and Layer 3 switches, and uninterruptible power supplies (UPSs).

Table 2. Breakdown of exposed devices using well-known SNMP read-only community strings, showing the percentage share of each device type within the total exposed sample ($N = 2403$).

Device Type	Observed Freq.	Relative Freq. (%)
Network Printer	1422	59.18%
Print Server	367	15.27%
Host	190	7.91%
Layer 3 Switch	113	4.70%
Layer 2 Switch	98	4.08%
IoT Sensor	45	1.87%
Label Printer	35	1.46%
Uninterruptible Power Supply (UPS)	27	1.12%
Storage Array	20	0.83%
Access Point	19	0.79%
Fibre Channel Switch	17	0.71%
Programmable Logic Controller (PLC)	9	0.37%
Firewall	6	0.25%
VoIP Phone	5	0.21%
Industrial Printer	4	0.17%
Load Balancer	4	0.17%
Network-Attached Storage (NAS)	4	0.17%
Tape Library	3	0.12%
VMware vCenter Server	3	0.12%
Embedded Module	2	0.08%
Environmental Monitoring Device (DCIM)	2	0.08%
Network USB hub	2	0.08%
Router	2	0.08%
3D Printer	1	0.04%
IP Camera	1	0.04%
Smart Building Server	1	0.04%
Wireless LAN Controller	1	0.04%
Total	2403	100.00%

Table 4 details the vendor breakdown of exposed network printers utilising well-known SNMP community strings. Regarding read-only operations (Table 4a), 1422 network printers are exposed, compared to 435 vulnerable to read-write operations (Table 4b). This results in a ratio of 3.27, which is highly consistent with the global ratio previously observed. As shown in Table 4a, six vendors display a relative frequency of read-only exposure of approximately 10% or higher. In addition, Table 4b reveals that only three vendors (Konica Minolta, Toshiba, and Sharp) show a prominent relative frequency when analysing read-write operations. Specifically, the network printer models most commonly identified with well-known SNMP read-write community strings are those belonging to the *bizhub i* series for Konica Minolta (https://manuals.konicaminolta.eu/bizhub-C360i-C300i-C250i-UD/EN/contents/id08_104516893.html, (accessed on 14 August 2026); https://manuals.konicaminolta.eu/bizhub-651i-551i-451i/EN/contents/WC_12_03_05.html, (accessed on 14 August 2026); https://manuals.konicaminolta.eu/bizhub-C4001i-C3301i/EN/contents/UT_05_05.html, (accessed on 14 August 2026)); the *e-STUDIO* series for Toshiba (https://business.toshiba.com/downloads/KB/f1Ulds/20520/eS2829A_TAG_EN_0002.pdf, see p. 48, (accessed on 14 August 2026)); and the *MX-Mx051* (<https://business.sharpusa.com/Portals/0/downloads/Manuals/Monochrome-Advanced-and-Essentials2-user.pdf>, see p. 740, (accessed on 14 August 2026)) and *BP-70C* (https://docs.aws.sharp.eu/Marketing/Operational_manuals/bp70c65_usr_02a_en.pdf, see p. 949, (accessed on 14 August 2026)) series for Sharp. For all of these series, official vendor manuals confirm that SNMP is enabled by default with both *public* and *private* community strings.

Table 3. Breakdown of exposed devices using well-known SNMP read–write community strings, showing the percentage share of each device type within the total exposed sample ($N = 645$).

Device Type	Observed Freq.	Relative Freq. (%)
Network Printer	435	67.44%
IoT Sensor	44	6.82%
Layer 3 Switch	28	4.34%
Layer 2 Switch	28	4.34%
Uninterruptible Power Supply (UPS)	22	3.41%
Label Printer	18	2.79%
Access Point	17	2.64%
Fibre Channel Switch	14	2.17%
Storage Array	10	1.55%
Print Server	9	1.40%
PLC	8	1.24%
Host	6	0.93%
Environmental Monitoring Device (DCIM)	2	0.31%
Firewall	2	0.31%
Wireless LAN Controller	1	0.16%
Network USB hub	1	0.16%
Total	645	100.00%

Table 4. Breakdown of exposed network printers.

Vendor	Observed Freq.	Relative Freq. (%)
(a) Well-known SNMP read–only community strings		
HP	295	20.75%
Konica Minolta	277	19.48%
Canon	228	16.03%
Kyocera	206	14.49%
Ricoh	136	9.56%
Toshiba	123	8.65%
Sharp	45	3.16%
Brother	44	3.09%
Xerox	33	2.32%
Lexmark	22	1.55%
Samsung	7	0.49%
Epson	2	0.14%
NRG	2	0.14%
Pantum	2	0.14%
Total	1422	100.00%
(b) Well-known SNMP read–write community strings		
Konica Minolta	239	54.94%
Toshiba	113	25.98%
Sharp	37	8.51%
Xerox	29	6.67%
HP	8	1.84%
Ricoh	4	0.92%
Epson	2	0.46%
Pantum	2	0.46%
NRG	1	0.23%
Total	435	100.00%

Table 5 shows the vendor breakdown of exposed print servers. Regarding read-only operations (Table 5a), 367 print servers are exposed, with HP representing the vast majority of cases, followed by Brother. As for Table 5b, it reveals a drastically lower exposure under read–write operations with only 9 cases. This asymmetry results in a read-only to read–write ratio of 40.77, deviating significantly from the global ratio previously observed.

Table 5. Breakdown of exposed print servers.

Vendor	Observed Freq.	Relative Freq. (%)
(a) Well-known SNMP read-only community strings		
HP	276	75.20%
Brother	87	23.71%
D-Link	4	1.09%
Total	367	100.00%
(b) Well-known SNMP read–write community strings		
HP	9	100.00%
Total	9	100.00%

Table 6 details the vendor breakdown of exposed hosts. With regard to read-only operations (Table 6a), 190 hosts are exposed, with *Linux* accounting for most cases, closely followed by *Windows*. Within this category, the identified devices primarily correspond to servers including, for instance, database and Active Directory (AD) servers. In addition, Table 6b reveals a drastically lower exposure under read–write operations with only 6 cases. This asymmetry results in a read-only to read–write ratio of 31.67, deviating significantly from the global ratio previously observed.

Table 6. Breakdown of exposed hosts.

Vendor	Observed Freq.	Relative Freq. (%)
(a) Well-known SNMP read-only community strings		
Linux	98	51.58%
Windows	90	47.37%
IBM i 7.4	1	0.53%
Unknown	1	0.53%
Total	190	100.00%
(b) Well-known SNMP read–write community strings		
Linux	4	66.67%
Windows	2	33.33%
Total	6	100.00%

Table 7 shows the vendor breakdown of exposed Layer 2 switches. Regarding read-only operations (Table 7a), 98 devices are exposed, with Cisco and HP appearing with a high relative frequency. Specifically, in the case of Cisco, the most frequently observed devices are *Cisco Catalyst 2960* switches, particularly the 2960-S and 2960-X series. In the case of HP, they belong to the *ProCurve* product family, mainly the 2530 series. Notably, while vendor documentation and tutorials indicate that several models within the *ProCurve* product family come with SNMP natively enabled for read-only operations, this default configuration does not apply to the aforementioned Cisco series.

Table 7. Breakdown of exposed Layer 2 switches.

Vendor	Observed Freq.	Relative Freq. (%)
(a) Well-known SNMP read-only community strings		
Cisco	45	45.92%
HP	27	27.55%
3Com	7	7.14%
Extreme Networks	7	7.14%
D-Link	5	5.10%
Aruba	4	4.08%
Enterasys Networks	1	1.02%
Netgear	1	1.02%
Unknown	1	1.02%
Total	98	100.00%
(b) Well-known SNMP read-write community strings		
Extreme Networks	7	25.00%
3Com	6	21.43%
Cisco	4	14.29%
D-Link	4	14.29%
HP	3	10.71%
MTL	3	10.71%
Unknown	1	3.57%
Total	28	100.00%

As for Table 7b, it reveals a much more uniform distribution across multiple vendors under read-write operations, totalling 28 cases. Notably, for this specific asset category, the read-only to read-write ratio stands at 3.50, closely aligning with the global ratio previously observed. It should be noted that this read-write capability is not enabled by default on any of the identified Layer 2 switches.

Table 8 details the vendor breakdown of exposed Layer 3 switches. Two specific vendors exhibit a substantially higher relative frequency than the rest in both scenarios, although they differ between cases. With regard to read-only operations (Table 8a), 113 devices are exposed, with Aruba and HP dominating the sample. Specifically, the most frequently observed models for Aruba belong to the 3500 series, whereas for HP they correspond to the 5400R series, both belonging to the broader *ProCurve* product lineage. In both cases, technical vendor documentation indicates that SNMP is natively enabled for read-only operations using the default *public* community string.

In addition, Table 8b reveals 28 cases under read-write operations, where Avaya and HP emerge as the prominent vendors. In this case, the most commonly identified models for Avaya are those from the 3500 series (for HP, once more the 5400 series). Again, this read-write capability is not enabled by default on any of the identified Layer 3 switches. Finally, the read-only to read-write ratio for this asset category stands at 4.04, which is highly consistent with the global ratio previously observed.

Table 9 shows exposed storage arrays. While Table 9a shows 20 read-only cases prominently featuring Dell, Table 9b reveals 10 read-write cases where Dell is entirely absent. Notably, this category's read-only to read-write ratio (2.00) falls below the global baseline, suggesting that exposed storage arrays, such as Hitachi and HP, tend to have both well-known community strings enabled simultaneously. It should be noted that, in this case, this is likely attributable to manual configuration by the network administrators, since none of the identified storage array models have SNMP enabled by default.

Table 8. Breakdown of exposed Layer 3 switches.

Vendor	Observed Freq.	Relative Freq. (%)
(a) Well-known SNMP read-only community strings		
Aruba	57	50.44%
HP	32	28.32%
Avaya	10	8.85%
Dell	4	3.54%
Cisco	3	2.65%
Extreme Networks	3	2.65%
Alcatel	1	0.88%
Netgear	1	0.88%
Nvidia	1	0.88%
Ubiquiti	1	0.88%
Total	113	100.00%
(b) Well-known SNMP read-write community strings		
Avaya	16	57.14%
HP	7	25.00%
Extreme Networks	2	7.14%
Cisco	1	3.57%
D-Link	1	3.57%
Netgear	1	3.57%
Total	28	100.00%

Table 9. Breakdown of exposed storage arrays.

Vendor	Observed Freq.	Relative Freq. (%)
(a) Well-known SNMP read-only community strings		
Dell	7	35.00%
HP	6	30.00%
Hitachi	3	15.00%
NetApp	3	15.00%
Netgear	1	5.00%
Total	20	100.00%
(b) Well-known SNMP read-write community strings		
HP	6	60.00%
Hitachi	4	40.00%
Total	10	100.00%

Table 10 details the vendor breakdown of exposed uninterruptible power supplies (UPSs). Regarding read-only operations (Table 10a), 27 devices are exposed, with APC and GenereX representing the vast majority of cases, specifically the APC AP9600 and the GenereX CS141 families. In addition, Table 10b reveals 22 cases under read-write operations, where the distribution remains heavily dominated by the same two manufacturers. Notably, this behaviour is even more exacerbated in this asset category, presenting a read-only to read-write ratio of 1.23.

This remarkably low ratio, dropping significantly below the global trend, strongly indicates that, when exposed, UPS devices almost systematically have both well-known community strings enabled simultaneously. In fact, both the APC AP9600 SmartSlot family (AP9606) (<https://www.manualsdir.com/manuals/47796/apc-ap9606.html?page=17>, see p. 13, (accessed on 14 August 2026)) and the GenereX CS141 (<https://www.genereX.de/>

media/pages/packages/documents/manuals/f65348d5b6-1783075459/manual_CS141_en.pdf, see p. 77, (accessed on 14 August 2026)) family are configured with the SNMP *public* and *private* community strings by default.

Table 10. Breakdown of exposed uninterruptible power supplies.

Vendor	Observed Freq.	Relative Freq. (%)
(a) Well-known SNMP read-only community strings		
APC	10	37.04%
Generex	9	33.33%
Riello	4	14.81%
Unknown	2	7.41%
PowerWalker	1	3.70%
UPS	1	3.70%
Total	27	100.00%
(b) Well-known SNMP read-write community strings		
APC	13	59.09%
Generex	6	27.27%
Riello	2	9.09%
PowerWalker	1	4.55%
Total	22	100.00%

5.4. Discovered CVEs, Misconfigurations and Technical Flaws

During the reconnaissance carried out across the networks of the 78 analysed organisations, twelve SNMP-related vulnerabilities were identified by the vulnerability scanner, as summarised in Table 11. It should be noted that these entries represent scanner-flagged vulnerabilities based on automated detection signatures, rather than manually exploited or independently verified flaws. Both the severity levels and descriptions for these vulnerabilities are referenced from the National Institute of Standards and Technology (NIST) database [34].

According to the Common Vulnerability Scoring System (CVSS), seven of these vulnerabilities present a medium severity level, whereas the remaining five are classified as high severity, with scores reaching up to 8.8. All are scored using CVSS v3.1, except for one medium severity (v3.0) and one high severity (v2.0) vulnerability.

The consequences of these security flaws vary significantly, ranging from causing a DoS condition on the SNMP service to achieving arbitrary command execution with root privileges. Furthermore, the threat posed by these CVEs is significantly increased by the widespread presence of well-known community strings documented in this study, as several of these flaws can be directly exploited.

Furthermore, based on the findings of the security assessments, a set of prevalent administrative misconfigurations and technical flaws was identified across several analysed organisations. These issues represent significant vectors of risk that should be strictly avoided within internal corporate environments:

- As evidenced by the results, numerous devices, most notably network printers, were deployed in their factory-default state, with the SNMP protocol natively active out of the box utilising default community strings (*public/private*).
- Tangentially related to the previous point and frequently discovered during SNMP scanning, many of these printers exposed web-based configuration interfaces. In several cases, they either relied on well-documented vendor default credentials or lacked authentication entirely, allowing unrestricted access to device settings.

- In addition to factory defaults, numerous network devices, predominantly Layer 2 and Layer 3 switches, were found with SNMP intentionally enabled and well-known community strings manually assigned by administrators, such as *public* or *private*. This blurs the distinction between default exposure and poor credential selection.
- Several devices accepted multiple read-only or read-write community strings. While this approach is not inherently insecure, it poses a severe risk when all configured strings correspond to common dictionary entries, such as the *Dict.txt* dictionary used by *Onesixtyone*. This flaw was predominantly observed within Layer 2 and Layer 3 switches, as well as storage arrays.
- Related to the previous point, and most critically, several storage arrays were identified as accepting any arbitrary value as a valid read-only or read-write community string. This behaviour indicates a fundamental authentication failure that completely bypasses access controls, leaving such devices entirely exposed to exploitation.

Table 11. Summary of discovered CVEs and associated severity.

CVE ID	Severity	Description
1999-0499	7.5 (high) [†]	NETBIOS shared information may be published through SNMP registry keys in Windows NT-based systems
2008-4309	7.5 (high)	Heap-based buffer overflow via crafted SNMP GET-BULK requests, causing a DoS (crash)
2018-18065	6.5 (medium) *	NULL pointer exception via crafted UDP packets, allowing authenticated remote attackers to cause a DoS (crash)
2020-15862	7.8 (high)	Improper privilege management via WRITE access to EXTEND MIB, allowing arbitrary command execution as root
2022-24805	8.8 (high)	Read-only users can trigger an out-of-bounds memory access via a buffer overflow in NET-SNMP-VACM-MIB INDEX handling
2022-24806	5.3 (medium)	Improper input validation when simultaneously SETing malformed OIDs in master and subagents, exploitable by read-write users
2022-24807	6.5 (medium)	Read-write users can trigger an out-of-bounds memory access via a malformed OID in a SET request to vacmAccessTable
2022-24808	6.5 (medium)	Read-write users can trigger a NULL pointer dereference via a malformed OID in a SET request to nsLogTable
2022-24809	6.5 (medium)	Read-only users can trigger a NULL pointer dereference via a malformed OID in a GET-NEXT request to nsVacmAccessTable
2022-24810	8.8 (high)	Read-write users can trigger a NULL pointer dereference via a malformed OID in a SET request to nsVacmAccessTable
2022-44792	6.5 (medium)	Write-access users can trigger a NULL pointer exception via a crafted UDP packet, causing a DoS (crash)
2022-44793	6.5 (medium)	Remote attackers can trigger a NULL pointer exception via a crafted UDP packet, causing a DoS (crash)

[†] CVSS version 2.0. * CVSS version 3.0. **Note:** Red indicates High severity vulnerabilities; orange indicates Medium severity vulnerabilities.

6. Discussion

The security flaws identified in this study extend beyond mere technical misconfigurations. This section examines potential underlying factors that may contribute to widespread SNMP exposure and proposes strategic changes in network administration and device deployment.

6.1. Shared Responsibility: Vendors and Suppliers

The results presented in this study suggest that unmodified factory default configurations constitute a key contributing factor in the widespread exposure of well-known SNMP community strings. Hardware vendors bear a foundational responsibility, as security-by-default principles require that protocols like SNMP must be consistently disabled by default across all vendors and devices. Furthermore, equipment suppliers and system installers, who oversee the physical installation and initial configuration for end customers, may also bear a significant share of the responsibility. These intermediaries might frequently focus on immediate operational availability rather than security hardening, deploying hardware in its factory state to facilitate a quicker installation. To mitigate this risk, supplier and system installer technicians could benefit from formal training regarding the protocols embedded within the devices they distribute, ensuring that unneeded services are systematically deactivated before completing the deployment.

6.2. Administrative Inconsistencies Stemming from Insufficient Training

A major concern is the administrative inconsistency identified within Layer 2 and Layer 3 switches, where network administrators intentionally enabled SNMP but manually configured weak community strings (even reusing conventional factory defaults like *public* or *private*). This practice may reflect a potential lack of understanding of SNMP behaviour. In many corporate environments, community strings appear to be treated merely as tokens to enable the protocol based on basic online tutorials, rather than the authentication credentials they fundamentally are. This situation might stem from limited awareness or specific training gaps, as network administrators may fail to treat community strings with the same relevance as corporate passwords. Leaving a core switch manageable by using a manually assigned weak community string is equivalent to exposing critical infrastructure through a default administrative password, granting attackers complete control over internal traffic routing.

6.3. The Underestimated Risks of Read-Only Operations

The gap between read-only (91%) and read-write (64%) exposure can lead to the misconception that the overall SNMP security posture is not critical. Network administrators might assume that read-only access is not dangerous because it prevents unauthorised configuration changes. However, during reconnaissance, read-only access provides attackers with a comprehensive network blueprint, exposing open ports, routing tables, active interfaces, and system descriptions. Furthermore, as shown by the active vulnerabilities identified in this study (such as CVE-2022-24805), authenticated read-only operations can lead to memory corruption or DoS conditions. This suggests that read-only exposure remains a high-severity risk vector.

6.4. Infrastructure Exposure by Organisation Size

The results reveal a more nuanced relationship between organisation size and infrastructure exposure than a simple linear trend. While read-only asset exposure increases steadily with size, read-write exposure and network device exposure instead rise progressively from S1 through S3 before declining in S4 (over 2000 employees).

This pattern suggests a potential non-linear relationship between organisation size and overall security posture. While not directly quantified in this study, a plausible interpretation is that small organisations (S1) benefit from reduced network complexity, which naturally limits exposure. As organisations expand into S2 and S3, network heterogeneity may increase faster than formal governance mechanisms, potentially explaining the peak in vulnerability observed at S3 (501–2000 employees). Conversely, the decline in exposure for the largest organisations (S4) could be associated with the deployment of dedicated administration teams and formalised policies that typically become operational at larger scales. Consequently, mid-to-large organisations appear to represent a vulnerable transition phase, having outgrown structural simplicity without necessarily having established enterprise-level governance.

6.5. The Role of Endpoints in Network Security

The high incidence of exposure among network printers, print servers, and hosts suggests that security efforts may not be fully focused on these devices. They are frequently misconfigured with factory default credentials and community strings, or even entirely unauthenticated web interfaces, making them accessible to external actors. Consequently, instead of being disregarded, these devices must be considered as active network elements with the same level of importance as the rest of the infrastructure. Otherwise, their exposure expands the attack surface, allowing attackers to leverage endpoints to compromise the organisational network.

Although the analysed organisations are geographically limited to Spain, this constraint does not significantly limit the generalizability of the findings. Spain belongs to the European Union and operates within its unified market for network equipment. Consequently, the hardware models, factory default configurations, and deployment practices observed in this study are representative of those adopted across all 27 EU member states, as well as partner nations bound by common commercial agreements. Furthermore, network administrator training across European countries follows similar academic and professional frameworks, reducing the likelihood of significant differences in security practices due to educational factors. Therefore, the factors identified as potentially contributing to SNMP exposure, including insecure default configurations, inadequate hardening practices, and insufficient awareness of SNMP security implications, might remain relevant beyond the analysed organisations and could be extended to similar environments in other countries.

7. Conclusions

This paper examines asset exposure resulting from the use of legacy SNMPv1 and SNMPv2c in private corporate networks across 78 Spanish organisations, using a grey-box penetration testing methodology. This approach uncovered severe misconfigurations, high-risk exposed devices, and known CVE-indexed vulnerabilities in live corporate networks, providing empirical evidence of their actual security posture beyond theoretical governance requirements. Although the geographical location of the organisations tested is limited, the outcomes and insights gathered may be extended to other countries. Consequently, these findings serve as a practical reference for network administrators seeking to enhance SNMP security. Additionally, this study provides administrative guidelines and deployment recommendations to mitigate these prevalent security flaws.

The empirical evaluation reveals widespread vulnerabilities across the evaluated organisations, with network infrastructure components—such as Layer 2 and Layer 3 switches—and endpoint assets being particularly affected. The analysis identified recurring technical flaws, including devices with SNMP enabled by default with well-known read-

only and read–write community strings and unauthenticated web management interfaces. Furthermore, several active CVEs were identified.

When interpreting these findings, several methodological limitations of the study should be considered. First, a purposive sampling strategy was employed rather than a random probability sampling method, as random selection in internal cybersecurity audits is often unfeasible due to authorisation and operational constraints. Although organisations of varying sizes and sectors were included, the non-random sampling introduces potential selection bias, restricting the analysis to descriptive statistics rather than population-wide inferential claims. Second, although the dataset covers 78 organisations, it is geographically limited to Spain. Nevertheless, the widespread deployment of standardised enterprise hardware, default vendor configurations, and common technical frameworks across European markets suggests that these findings could also be observed in other countries. Finally, due to internal routing policies and administrative boundaries, complete coverage of all organisation-wide subnets was not achievable. This partial coverage represents a key methodological limitation. However, because all reported vulnerabilities represent confirmed findings with zero false positives, this restriction renders our results in a conservative lower bound of total exposure.

Future work aims to expand this research to a broader sample of organisations internationally, collaborating with external entities and researchers to assess SNMP exposure across different sectors and countries. Furthermore, a longitudinal study will be considered to track the evolution of this exposure over time within the specific organisations analysed in this work.

Author Contributions: Conceptualization, P.N., D.Á.; Design, P.N., D.Á., F.G.B., A.F., J.C.G.; Research, P.N., F.G.B.; Methodology, P.N., D.Á., F.G.B., A.F., J.C.G.; Resources, P.N., D.Á., F.G.B.; Investigation, P.N., D.Á., A.F.; Data Curation, D.Á., A.F.; Writing, P.N., D.Á., F.G.B., A.F., J.C.G.; Supervision, P.N., J.C.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Informed consent for participation was obtained from all organisations involved in the study.

Data Availability Statement: Due to privacy and ethical restrictions, as well as the non-disclosure agreements signed with the participating organisations, the datasets generated and analysed during this study are not publicly available to ensure the anonymity and security of the evaluated corporate networks.

Conflicts of Interest: The authors declare no potential conflicts of interests.

References

1. Stallings, W. SNMP and SNMPv2: The infrastructure for network management. *IEEE Commun. Mag.* **1998**, *36*, 37–43. [\[CrossRef\]](#)
2. Li, T.; Yang, C.; Wang, Y.; Cai, L.; Anpalagan, A.; Han, Z. A survey on network management for xANET: Evolution, challenges, future directions. *IEEE Commun. Surv. Tutor.* **2025**, *28*, 1209–1247. [\[CrossRef\]](#)
3. Ferdoush, T.E.; Hasan, M. Technical analysis and design of a secure campus area network and monitoring system. In *Proceedings of the 2023 International Conference on Information and Communication Technology for Sustainable Development (ICICT4SD)*, Dhaka, Bangladesh; IEEE: Piscataway, NJ, USA, 2023; pp. 17–20. [\[CrossRef\]](#)
4. Kant, D.; Creutzburg, R.; Johannsen, A. Investigation of risks for critical infrastructures due to the exposure of SCADA systems and industrial controls on the Internet based on the search engine Shodan. *Electron. Imaging* **2020**, *32*, 253.1–253.15. [\[CrossRef\]](#)
5. Murthy, A.; Asghar, M.R.; Tu, W. A lightweight intrusion detection for Internet of Things-based smart buildings. *Secur. Priv.* **2024**, *7*, e386. [\[CrossRef\]](#)
6. Liu, H.; Huo, Y. Analysis on the security of MIB objects defined in network device. In *Proceedings of the 2016 2nd IEEE International Conference on Computer and Communications (ICCC)*, Chengdu, China; IEEE: Piscataway, NJ, USA, 2016; pp. 1067–1070. [\[CrossRef\]](#)

7. Dietrich, C.; Krombholz, K.; Borgolte, K.; Fiebig, T. Investigating system operators' perspective on security misconfigurations. In *CCS '18: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, New York, NY, USA*; ACM: New York, NY, USA, 2018; pp. 1272–1289. [\[CrossRef\]](#)
8. Tiefenau, C.; Häring, M.; Krombholz, K.; Von Zezschwitz, E. Security, availability, and multiple information sources: Exploring update behavior of system administrators. In *SOUPS'20: Proceedings of the Sixteenth USENIX Conference on Usable Privacy and Security, USA*; USENIX Association: Berkeley, CA, USA, 2020; pp. 239–258. Available online: <https://www.usenix.org/conference/soups2020/presentation/tiefenau> (accessed on 17 July 2026).
9. Loureiro, S. Security misconfigurations and how to prevent them. *Netw. Secur.* **2021**, *2021*, 13–16. [\[CrossRef\]](#)
10. May, R.; Biermann, C.; Krüger, J.; Leich, T. Asking security practitioners: Did you find the vulnerable (mis)configuration? In *VaMoS '25: Proceedings of the 19th International Working Conference on Variability Modelling of Software-Intensive Systems, New York, NY, USA*; ACM: New York, NY, USA, 2025; pp. 30–39. [\[CrossRef\]](#)
11. Álvarez, D.; Nuño, P.; Bulnes, F.G.; Granda, J.C. Network security misconfigurations in Spanish organisations: An empirical study. *Comput. Netw.* **2026**, *268*, 112631. [\[CrossRef\]](#)
12. Onesixtyone Fast SNMP Scanner. Available online: <https://github.com/trailofbits/onesixtyone> (accessed on 17 July 2026).
13. Czyz, J.; Luckie, M.; Allman, M.; Bailey, M. Don't forget to lock the back door! A characterization of IPv6 network security policy. In *Proceedings of the Network and Distributed Systems Security (NDSS), San Diego, CA, USA*; Internet Society: Reston, VA, USA, 2016; pp. 1–15. [\[CrossRef\]](#)
14. Bjerre, A.; Westh, A.P.; Villefrance, E.; Haque, A.S.M.F.A.; Andersen, J.B.; Helgogaard, L.K.; Anagnostopoulos, M. A wide network scanning for discovery of UDP-based reflectors in the Nordic countries. In *Proceedings of the Secure IT Systems, Reykjavic, Iceland*; Springer: Cham, Switzerland, 2022; pp. 176–193. [\[CrossRef\]](#)
15. Samtani, S.; Yu, S.; Zhu, H.; Patton, M.; Matherly, J.; Chen, H. Identifying SCADA systems and their vulnerabilities on the Internet of Things: A text-mining approach. *IEEE Intell. Syst.* **2018**, *33*, 63–73. [\[CrossRef\]](#)
16. Deng, Q.; Pu, J.; Tan, Z.; Qian, Z.; Krishnamurthy, S.V. Beyond the horizon: Uncovering hosts and services behind misconfigured firewalls. In *Proceedings of the 2025 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA*; IEEE: Piscataway, NJ, USA, 2025; pp. 1770–1788. [\[CrossRef\]](#)
17. Boyko, A.; Varkentin, V.; Polyakova, T. Advantages and disadvantages of the data collection's method using SNMP. In *Proceedings of the 2019 International Multi-Conference on Industrial Engineering and Modern Technologies (FarEastCon), Vladivostok, Primorsky Krai, Russia*; IEEE: Piscataway, NJ, USA, 2019; pp. 1–5. [\[CrossRef\]](#)
18. Wang, Z.; Zhang, Y.; Liu, Q. RPFuzzer: A framework for discovering router protocols vulnerabilities based on fuzzing. *KSII Trans. Internet Inf. Syst.* **2013**, *7*, 1989–2009. [\[CrossRef\]](#)
19. Feng, X.; Wu, J.; Wang, K.; Li, J.; Wang, M. Security analysis of Simple Network Management Protocol based IEEE P21451 Internet of Things. In *Proceedings of the 2017 IEEE 15th Intl Conf on Dependable, Autonomic and Secure Computing, 15th Intl Conf on Pervasive Intelligence and Computing, 3rd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress (DASC/PiCom/DataCom/CyberSciTech), Orlando, FL, USA*; IEEE: Piscataway, NJ, USA, 2017; pp. 326–330. [\[CrossRef\]](#)
20. Fan, H.; Dong, Y.; Yu, M.; Tung, L. Security threats against the communication networks for traffic control systems. In *Proceedings of the 2013 IEEE International Conference on Systems, Man, and Cybernetics, Manchester, UK*; IEEE: Piscataway, NJ, USA, 2013; pp. 4783–4788. [\[CrossRef\]](#)
21. Gondim, J.J.; de Oliveira Albuquerque, R.; Sandoval Orozco, A.L. Mirror saturation in amplified reflection distributed denial of service: A case of study using SNMP, SSDP, NTP and DNS protocols. *Future Gener. Comput. Syst.* **2020**, *108*, 68–81. [\[CrossRef\]](#)
22. Lin, S. Research on security of UDP reflection attack in civil aviation information network based on SDN. In *Proceedings of the 2020 IEEE 2nd International Conference on Civil Aviation Safety and Information Technology (ICCASIT), Weihai, China*; IEEE: Piscataway, NJ, USA, 2020; pp. 210–216. [\[CrossRef\]](#)
23. Albakour, T.; Gasser, O.; Beverly, R.; Smaragdakis, G. Third time's not a charm: Exploiting SNMPv3 for router fingerprinting. In *IMC '21: Proceedings of the 21st ACM Internet Measurement Conference, New York, NY, USA*; ACM: New York, NY, USA, 2021; pp. 150–164. [\[CrossRef\]](#)
24. Albakour, T.; Gasser, O.; Beverly, R.; Smaragdakis, G. Illuminating router vendor diversity within providers and along network paths. In *IMC '23: Proceedings of the 2023 ACM on Internet Measurement Conference, New York, NY, USA*; ACM: New York, NY, USA, 2023; pp. 89–103. [\[CrossRef\]](#)
25. Otrok, H.; Mourad, A.; Debbabi, M.; Assi, C. Improving the security of SNMP in wireless networks. In *Proceedings of the 2005 International Conference on Wireless Networks, Communications and Mobile Computing, Maui, HI, USA*; IEEE: Piscataway, NJ, USA, 2005; Volume 1, pp. 198–202. [\[CrossRef\]](#)
26. Madhusudhan, R.; Shashidhara. Cross channel scripting (XCS) attacks in web applications: Detection and mitigation approaches. In *Proceedings of the 2018 2nd Cyber Security in Networking Conference (CSNet), Paris, France*; IEEE: Piscataway, NJ, USA, 2018; pp. 1–3. [\[CrossRef\]](#)

27. Khan, H.M.A.; Sarwar, N.; Hamdi, M.; Inayat, U. Trusted simple network management protocol for integrity evaluation in distributed environment. *J. Netw. Syst. Manag.* **2026**, *34*, 48. [[CrossRef](#)]
28. Huq, N.; Hilt, S.; Hellberg, N. *US Cities Exposed*; Technical Report; TrendMicro: Tokyo, Japan, 2017. Available online: <https://infosecuritymagazine.be/files/cd42a564c04be8ad28dd2f08c5d45933.pdf> (accessed on 17 July 2026).
29. Hellberg, N.; Vosseler, R. *Western European Cities Exposed*; Technical Report; TrendMicro: Tokyo, Japan, 2017. Available online: https://techfromthenet.it/wp-content/uploads/2017/12/documents.trendmicro.com_assets_wp_wp-western-europe-cities-exposed.pdf (accessed on 17 July 2026).
30. Hellberg, N.; Vosseler, R. *German Cities Exposed*; Technical Report; TrendMicro: Tokyo, Japan, 2017. Available online: <https://documents.trendmicro.com/assets/wp/wp-german-cities-exposed.pdf> (accessed on 17 July 2026).
31. Rajasekar, V.; Rajkumar, S. A study on Internet of Things devices vulnerabilities using Shodan. *Int. J. Comput.* **2023**, *22*, 149–158. [[CrossRef](#)]
32. Kepka, I. Analysis of SNMP through Censys datasets. In *Proceedings of the 41st Twente Student Conference on IT (TSC 41)*; University of Twente: Enschede, The Netherlands, 2024; pp. 1–9. Available online: <https://drive.google.com/file/d/1a5kr-RbyaGWWn-4tWugD4aDqsCYZRvJv/view> (accessed on 17 July 2026).
33. Bhardwaj, A.; Sapra, V.; Sapra, L. Evading firewalls & enumerate SNMP using advanced NMAP techniques. In *Proceedings of the 2023 3rd Asian Conference on Innovation in Technology (ASIANCON)*, Pune, India; IEEE: Piscataway, NJ, USA, 2023; pp. 1–6. [[CrossRef](#)]
34. National Vulnerability Database (NVD). Available online: <https://nvd.nist.gov/> (accessed on 17 July 2026).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.