Low-Rate Distributed Denial of Service across Sectors and Benchmark Datasets: A Comparative Survey

[1] Ms. Achsah Susan Mathew, [2] Dr. Veena R, [3] Dr. Hanumanthappa M

[1] Research Scholar, Department of Computer Science, Bangalore University, Bengaluru, Karnataka
[2] Professor, Department of Computer Science, DBIMSCA, Bangalore, Karnataka,
[3] Senior Professor, Department of Computer Science, Bangalore University, Bengaluru, Karnataka, India.

Abstract: Distributed Denial of Service (DDoS) attacks are no longer dominated by huge volumetric floods. A quieter but equally damaging class of attack the low-rate DDoS sends traffic at rates so low that most threshold-based defenders never raise an alarm. Slowloris, R.U.D.Y(R-U-Dead-Yet), Slow-Read and the classic Shrew attack all comes under the same family. They are especially dangerous in sectors where the underlying traffic is itself light and predictable, such as healthcare with its IoMT devices, banking portals at off-peak hours and small cloud tenants. This survey brings three things together. First, this describes the major sectors targeted by DDoS attacks, for each, the specific low-rate variants that show up most often. Second, it reviews the public benchmark datasets used in DDoS detection research and score each one on how well it covers low-rate variants an area where many otherwise popular datasets are surprisingly weak. Third, propose a sector-to-dataset mapping that takes low-rate coverage into account, so that a researcher choosing a dataset for low-rate DDoS work in a particular sector is not left training on volumetric-only data. The result is a single short reference for the first design decisions in a low-rate DDoS detection study.

Keywords **- Low-rate DDoS, Slowloris, Shrew attack, R.U.D.Y., intrusion detection, benchmark datasets, sector analysis, CICDDoS2019, CIC-IoMT2024, healthcare cybersecurity.**

1. INTRODUCTION

When people picture a DDoS attack, they usually imagine a flood millions of packets per second, traffic graphs spiking to the ceiling, defenders scrambling to absorb the load. That picture is true for one class of attack but not all of them. A second class, called low-rate DDoS, takes the opposite approach. Instead of trying to overwhelm the target with volume, it sends traffic that looks almost normal. It just sends it in a way that wastes server resources slowly, hour after hour [1], [22].

Three names keep coming up when people study low-rate DDoS. The first is Shrew attack. Kuzmanovic and Knightly came up with it back in 2003 [23]. It sends short bursts at just the right time. TCP gets fooled into thinking the network is jammed and slows itself down. Slowloris works in a totally different way. RSnake put it out as a tool in 2009 [24]. It opens hundreds of HTTP connections to a target server, then keeps each one alive by sending one byte at a time. R.U.D.Y., short for (R-U-Dead-Yet), is basically Slowloris but for HTTP POST bodies instead of GETs. None of them needs much bandwidth. That's why simple defenders almost never catch them.

For a researcher, this raises a question that the broader DDoS literature does not always answer clearly. If our goal is to detect low-rate DDoS, do the standard sector statistics and benchmark datasets actually serve us? Some sectors are more exposed to low-rate variants than others. And many of the most popular datasets including CICDDoS2019, which is the default choice for a lot of DDoS papers are weighted heavily toward volumetric reflection attacks and contain very little low-rate traffic [12], [25].
This paper answers three connected questions. (i) Which sectors face the highest exposure to low-rate DDoS, and why? (ii) Which public datasets actually contain enough low-rate samples to train a detection model? (iii) Given a sector, which dataset is the right one to pick? It shows dataset specifications and a sector-to-dataset mapping together in one place, and added a dedicated table that scores every dataset on its low-rate coverage.

2. LITERATURE SELECTION APPROACH

Followed three steps. First, collected attack statistics by sector from public reports published by Cloudflare, StormWall, NETSCOUT, Akamai and MazeBolt between 2023 and 2025 [1]-[5], and added 2024 - 2026 academic surveys from IEEE, Elsevier and Springer [6]-[10].
Secondly, collected the datasets cited most often in recent (2022–2026) DDoS detection papers from IEEE Xplore, ACM Digital Library and ScienceDirect. For each dataset we recorded the source URL, year, size, number of features, attack labels and importantly whether it contains labelled low-rate attack samples [11]-[21]. Third, mappings: a sector-to-dataset table that recommends datasets per sector, and a separate low-rate-coverage table that scores each dataset on its usefulness for low-rate DDoS research specifically.

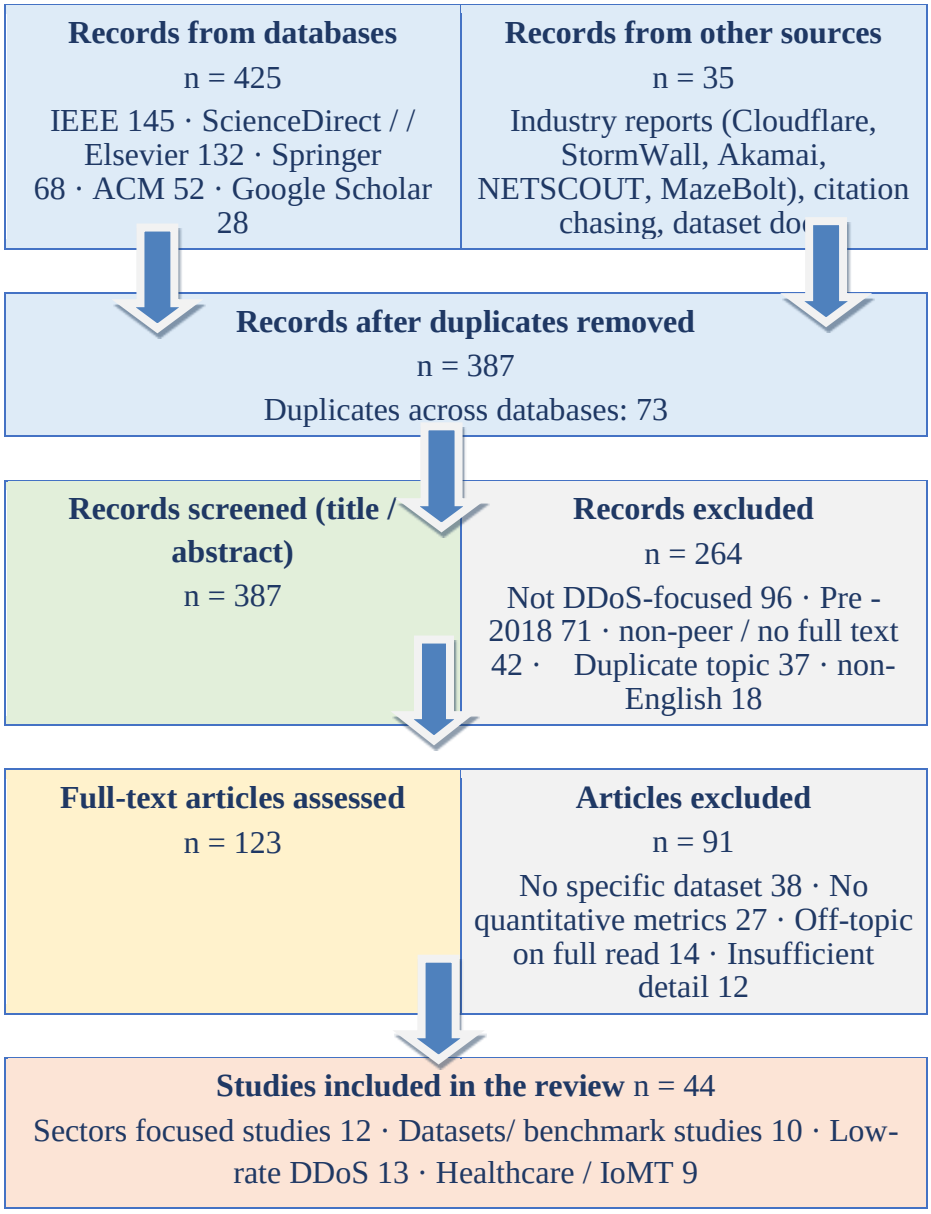

Fig 1. PRISMA-style flow of records through identification, screening, eligibility and inclusion

3. UNDERSTANDING LOW-RATE DDOS ATTACKS

Low-rate DDoS is best understood by contrast. A volumetric flood tries to be obvious: it pushes as much traffic at the target as possible, hoping to saturate either the bandwidth or the server's connection table. A low-rate attack does the opposite. It picks a specific server resource a TCP retransmission timer, an HTTP connection slot, an application thread pool and slowly exhausts it while keeping the overall bandwidth low enough to slip under most defender's radar.

Fig. 2 shows the traffic-rate difference. A volumetric flood produces a tall spike that any threshold-based defender catches in seconds. A low-rate attack produces a flat line that looks almost identical to baseline traffic, but it slowly degrades the server's health over many minutes or hours.

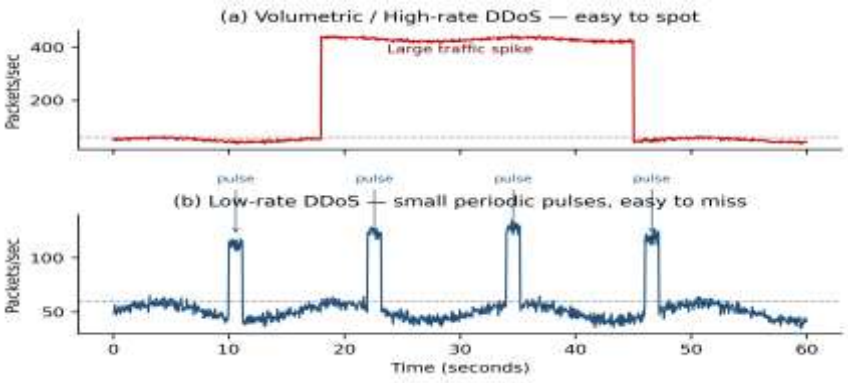

Fig 2. Traffic-rate comparison: volumetric DDoS (left) versus low-rate DDoS (right). Low-rate traffic stays close to the baseline, which is why simple threshold-based defenders miss it.

3.1 DDOS Attacks across the OSI layers

To understand where low-rate DDoS attacks fit in, first map the common DDoS attacks to the seven layers of the OSI reference model. Figure 3 shows this mapping. Most loud volumetric attacks operate at the Network (Layer 3) and Transport (Layer 4) layers, while the more subtle low-rate and slow attacks usually target the Application layer (Layer 7) [33], [34].

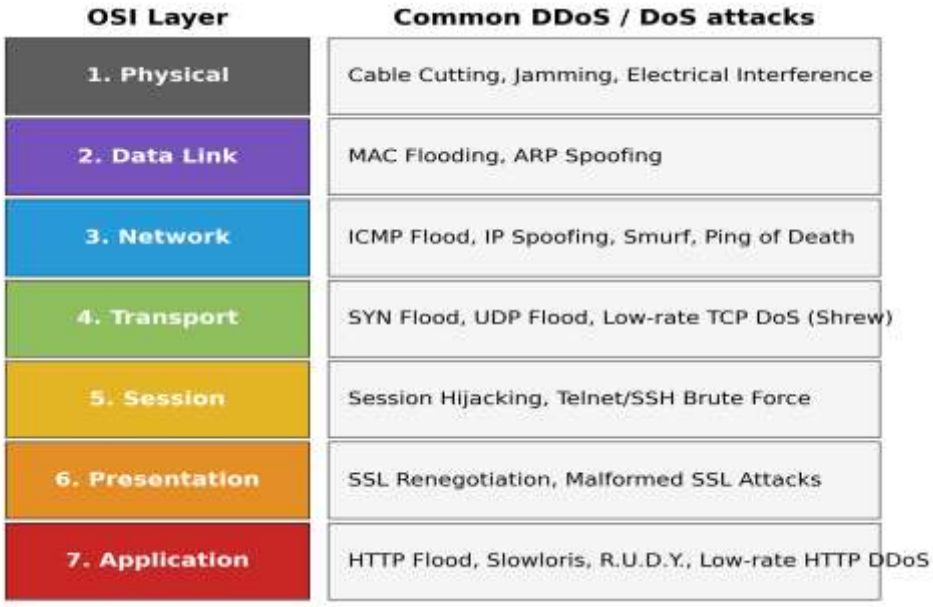

Fig 3. Mapping of common DDoS and DoS attacks to the seven layers of the OSI reference model.

3.1.1 Why low-rate is hard to detect

Three properties make low-rate detection difficult. The traffic is below typical alert thresholds, so volume-based detectors stay silent. The packets are usually well-formed (a valid HTTP GET with proper headers, or a normal TCP segment), Kumar et al. paper explicitly argues that today's commercially available intrusion detection systems are signature-based that are not capable of detecting unknown attacks so signature-based detectors find nothing to match [45]. And the attack is distributed across many sources, so per-source rate limits do not trigger either [22], [26].

3.1.2 Common low-rate attack variants

Fig. 4 groups the most common variants into two layers. Transport-layer variants (the Shrew family) target TCP's congestion-control mechanism. Various Application-layer attacks (Slowloris, R.U.D.Y., Slow Read) target web servers. SIP-targeted variants attack VoIP signaling and are increasingly seen in telecom-sector incidents [27].

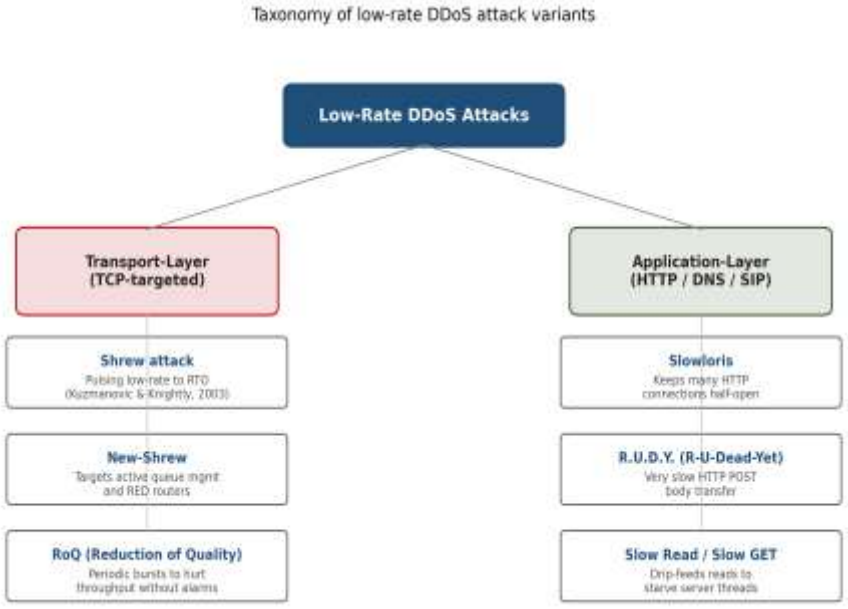

Fig. 4. Taxonomy of low-rate DDoS attack variants by target layer.

3.1.3 Why low-rate attacks are growing

Few trends are putting low-rate attacks upward in recent reports. First, cloud auto-scaling means a low-rate attack can directly inflate the victim's bill (the Economic Denial of Sustainability problem) [6]. Second, IoMT and IoT devices typically generate light, predictable traffic, so even a small disturbance is enough to disrupt them and a small disturbance is exactly what a low-rate attack

looks like [9], [25]. Third, defenders have gotten very good at filtering volumetric attacks, which pushes attackers toward quieter techniques [4].

4. SECTORS TARGETED BY DDOS ATTACKS

Fig. 5 shows the share of global DDoS incidents across most-targeted sectors in 2024 - 2025 industry reports

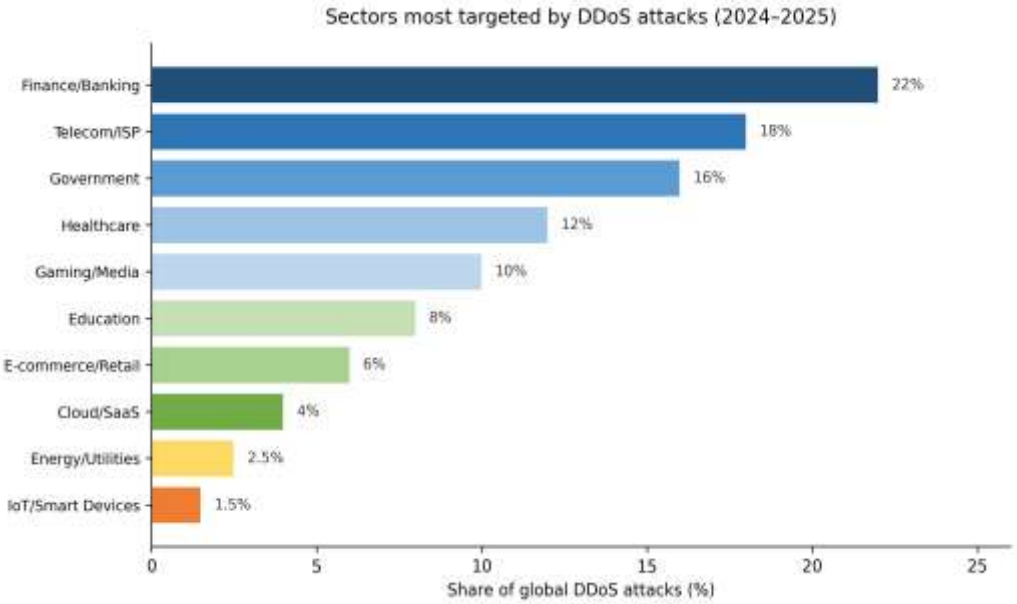

Fig 5. The ten sectors hit most often by DDoS attacks worldwide, with each sector's share of all incidents [1], [2], [3].

Six sectors matter most for low-rate DDoS research. Each one is covered below: Finance and Banking, Telecommunications and ISPs, Healthcare and IoMT, E-Commerce and Retail, Energy and Utilities, and IoT / Smart Devices.

4.1 Financial Services and Banking

Banks get hit by two kinds of DDoS. The loud, obvious kind and the quiet, sneaky kind. The obvious ones are the easy ones to spot. HTTP floods. DNS amplification. The bandwidth graph jumps to the ceiling, pagers start buzzing in the SOC, and the security team locks things down inside minutes. The sneaky kind is harder. Slowloris is the classic example. It barely sends any traffic at all. What it does is open up hundreds of half-finished HTTP connections to a bank's login page or transfer page, then sit on them. Bandwidth on the dashboard looks fine. But the page itself starts crawling. Customers can't log in. Transfers fail. Nobody figures out something is wrong until the support line gets flooded with angry callers. In 2024, someone got even cleverer. Reports say attackers started mixing the two kinds on purpose. A noisy flood goes out the front door, the security team chases it, and meanwhile a quiet slow attack walks in the back and does the actual damage [3].

4.2 Telecommunications and Internet Service Providers

Telecom is the backbone of Internet. That's what makes them such an obvious target for the big volumetric attacks. The kind that can knock out connectivity for an entire region for hours. The newer worry is quieter. Attackers have started going after the less visible parts of the network. Mainly SIP signaling and the way VoIP calls get routed. The trick is simple. Flood SIP REGISTER and INVITE requests at a low enough rate that bandwidth alerts don't trip. But it's still enough to clog the call-routing pipeline. VoIP service starts dropping calls. New calls get refused [27]. Bandwidth barely moves on the NOC dashboards. So, the ops team has no idea anything is wrong until customers start flooding the helpline saying their calls won't go through. Carriers also see a steady stream of reflection and amplification attacks. The attacker basically borrows the carrier's huge outbound pipe and points it at somebody else's server [1], [3].

4.3 Healthcare Services and IoMT

Healthcare is where low-rate DDoS does the most damage for the least amount of effort. Most medical devices on a hospital network don't send much data to begin with. A glucose monitor uploads a reading every few minutes. A heart-rate sensor sends a tiny packet every few seconds. That's it. That's the entire traffic pattern. So, when something doesn't fit that pattern, even slightly, the device starts behaving strangely. And nobody on the IT team can figure out why at first [9]. Things get worse a layer up. The IoMT security reviews from the past two years keep pointing at the same two problems. Slow-rate HTTP attacks chewing on telemedicine portals. Slow-read attacks slowly starving EMR system connections. Both kinds slide right past the dashboards [8], [25]. Two cases from the past couple of years really drive the point home. Change Healthcare in 2024. AZ Monica hospital in Belgium in 2026. Both attacks were slow and quiet. Both stayed undetected way longer than anyone expected [4].

4.4 Energy, Utilities and Critical Infrastructure

The energy and utilities sector reports almost no DDoS incidents publicly. But that doesn't mean much. Operators in this space tend to keep cyber problems under wraps. Some of it is about safety. Some of it is about how regulators handle disclosure. The few incidents that do leak out are usually slow-rate attacks against SCADA web interfaces and HMI portals. These are the screens engineers stare at when they are monitoring physical equipment in real time [5]. The bandwidth these systems use is already tiny. So, a low-rate attack doesn't need much traffic to cause real trouble. A small slowdown on an HMI dashboard is enough to delay

an operator's reaction to a real-world problem. And that delay is exactly the gap an attacker is hoping for. Life-safety stakes plus quiet attack patterns. That's what makes this sector worth worrying about. Low public incident counts don't change that.

4.5 E-Commerce and Retail

E-commerce sites have one big thing going against them. The traffic spikes are predictable Black Friday, Diwali. The big festival weeks. Year-end clearance windows. Attackers know exactly when the site's going to be busy. They also know exactly when a few seconds of downtime turns into thousands of orders walking away. The usual playbook is a mix. HTTP floods to slow the checkout flow. Slow-rate attacks aimed at the login and cart pages. And usually, a layer of bots pretending to be real shoppers while quietly scraping pricing data, locking out genuine customers in the process [2]. The result is everything an online retailer dreads. Carts abandoned at the payment step. Angry posts piling up on social media within minutes. Revenue draining away in real time. A fifteen-minute outage in the middle of a flash sale can cost more than a year's fraud-prevention budget.

4.6 Internet of Things and Smart Devices

When it comes to IoT and DDoS, things don't quite work as expected. Usually, IoT devices are the ones doing the attacking, not the ones being attacked. For instance, cameras, routers, and smart appliances can be compromised and used to create huge botnets like Mirai, Mozi, and the newer Aisuru family. These botnets then launch massive attacks on other areas [10], [17]. But, there's another side to IoT it can also be the victim. The gateways and cloud platforms that collect data from millions of devices are now being targeted with slow-rate attacks. The goal of these attacks is to disrupt the flow of data. If an attack is successful, it can cause big problems. For example, if a smart-city gateway is attacked, traffic signals might not get the latest information, and that can cause chaos. In a smart home, it might mean that the lights get stuck on. In an industrial plant, it can be even more serious machinery might not get the data it needs to work properly, so it's basically operating blind. This can have serious consequences, and it's something that needs to be taken seriously. Slow-rate attacks might not be as flashy as some other types of attacks, but they can still cause a lot of damage. So, it's essential to make sure that IoT devices and the systems that support them are secure and can withstand these types of attacks.

5. SECTORS TARGETED BY DDOS ATTACKS

Table I gives a side-by-side comparison of all ten sectors. It now includes a column for low-rate exposure Low / Medium / High so it will give a glance where low-rate DDoS is most concerning.

TABLE I. SECTOR-LEVEL COMPARISON WITH LOW-RATE DDOS EXPOSURE

Sector	Common low-rate variants	Share (%)	Low-rate exposure	Most damaging outcome
Finance / Banking	Slowloris, R.U.D.Y. on portals	22	High	Stalled transactions
Telecom / ISP	SIP low-rate, Shrew	18	Medium	Cascading customer outage
Healthcare / IoMT	Slow-rate IoMT, Slow Read EMR	12	High	Patient safety risk
E-commerce / Retail	Slowloris + bot scraping	6	Medium	Lost sales on peak days
Energy / Utilities	Slow HMI / SCADA portal	2.5	High	Critical service blackout
IoT / Smart Devices	Slow telemetry disruption	1.5	Low	Lost telemetry

6. BENCHMARK DATASETS FOR DDOS DETECTION RESEARCH

Picking a dataset is one of the most important early decisions in a DDoS detection project. For low-rate work the question is even more delicate, because the dataset must contain enough labelled low-rate samples to train and validate a model. In this section we summarize the public datasets cited most often in 2022 - 2026 DDoS detection papers, in roughly chronological order. Fig. 6 places them on a timeline so the reader can see how the field has shifted from generic IT traffic toward IoT and IoMT traffic in recent years.

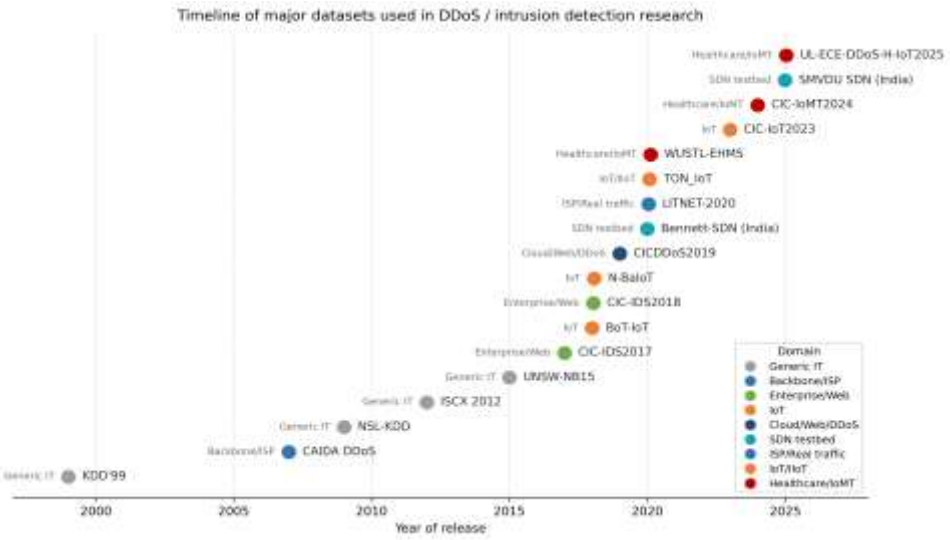

Fig 6. Timeline of major benchmark datasets used in DDoS / intrusion detection research, grouped by domain.

6.1 KDD'99 and NSL-KDD

Both contain DoS labels (mostly Neptune SYN-flood style) but no labelled low-rate variants. Useful only as a quick sanity baseline [11].

6.2 CAIDA DDoS Attack 2007

Real backbone traffic from one DDoS event. Mostly volumetric. Useful for ISP-style validation but does not help low-rate research directly [13].

6.3 UNSW- NB15

Includes a small set of DoS scenarios. Some application-layer slow-style attacks but not explicitly labelled as low-rate [14].

6.4 CIC-IDS2017 and CIC-IDS2018

These two datasets, released by the Canadian Institute for Cybersecurity, explicitly include Slowloris, SlowHTTPTest and Slow Read attacks alongside enterprise traffic. For low-rate DDoS research, CIC-IDS2017 and CIC-IDS2018 are arguably the strongest options in the entire list [15].

6.5 CICDDoS2019

The most cited DDoS-specific dataset but ironically the weakest for low-rate work. Its 13 attack classes are almost entirely reflective/volumetric (DNS, NTP, NetBIOS, LDAP, SSDP and so on). Researchers focused on low-rate DDoS often have to augment CICDDoS2019 with samples generated from Slowloris or hping3, as discussed in [12], [25].

6.6 BoT-IoT, TON_IoT and N-BaIoT

BoT-IoT and TON_IoT cover IoT/IIoT scenarios with mostly volumetric attacks. N-BaIoT focuses on Mirai- and Bashlite-style attacks. Limited low-rate content overall, although TON_IoT includes some application-layer DoS scenarios [16], [17].

6.7 CIC-IoT2023 and CIC-IoMT2024

CIC-IoT2023 includes 33 attack types in a smart-home testbed with 105 devices, of which a small set are HTTP slow-rate variants. CIC-IoMT2024 extends the same idea to 40 medical devices (25 real, 15 simulated) with 18 attacks across Wi-Fi, MQTT and Bluetooth, grouped into five categories: DDoS, DoS, Recon, MQTT and spoofing [18], [19]. An important caveat for low-rate research is that the DoS and DDoS samples in CIC-IoMT2024 are predominantly volumetric UDP, ICMP, SYN and MQTT connect floods make up the bulk of the data and the dataset does not include labelled slow-rate variants such as Slowloris or SlowITe. CIC-IoMT2024 is therefore the most relevant dataset for general IoMT security research, but for healthcare-focused low-rate DDoS work it needs to be augmented with low-rate samples (for example, from CIC-IDS2017 or by generating SlowITe-style traffic against an MQTT broker [30]).

6.8 WUSTL-EHMS-2020 and LITNET-2020

WUSTL-EHMS adds patient biometric features to network features unusual and valuable for healthcare studies but its attack labels are spoofing and injection, not low-rate DDoS [20]. LITNET-2020 provides ten months of real Lithuanian academic-network traffic across 12 attack categories, including some application-layer slow attacks [21].

6.9 INDIAN DATASETS: BENNETT-SDN, SMVDU

These publicly known DDoS datasets have been released by Indian institutions. The Bennett-SDN dataset by Ahuja, Singal and Mukhopadhyay (Bennett University, Greater Noida, 2020) is the most widely cited; it uses ten Mininet topologies with a Ryu SDN controller and contains TCP SYN flood, UDP flood and ICMP flood traffic across 23 flow features [35]. The SMVDU SDN TCP-SYN dataset by Kumar and Gupta (Shri Mata Vaishno Devi University, Katra, 2025) is narrower in scope, covering only the TCP-SYN flood scenario in an SDN environment [36].

6.10 UL-ECE-DDOS-H-IOT2025 (HEALTHCARE IOT, FREQUENCY-BASED)

Released by Akhi, Eising and Dhirani from the University of Limerick in 2025, this is the most recent and most directly relevant dataset for our research focus [37]. It contains two labelled CSV sub-datasets generated using the Cooja simulator for MQTT traffic and the ns-3 simulator for UDP traffic, modelling H-IoT devices such as body-temperature, oxygen-saturation and heart-rate sensors in a 5G-based healthcare environment. The attack model is interesting: normal H-IoT nodes transmit at 60-second intervals, while DDoS-affected nodes publish at 20-second intervals. This 3x frequency anomaly is not a classical flood it is much closer to a moderate- or low-rate behavioural attack, and the dataset is therefore valuable as a complement to CIC-IoMT2024 (which provides realistic IoMT device context) when training models that need to detect subtler attack patterns in healthcare environments.

7. BENCHMARK DATASETS FOR DDOS DETECTION RESEARCH

Table II summarizes the practical details of each dataset. Compared to general-purpose surveys, final column that flags whether the dataset contains labelled low-rate DDoS samples (Yes / Partial / No).

TABLE II. SPECIFICATIONS OF PUBLIC DATASETS WITH LOW-RATE COVERAGE

Dataset	Year	Origin	Domain	Size	Features	Attack labels	Low-rate samples?
CAIDA DDoS	2007	USA	Backbone / ISP	~21 GB	PCAP	1 real DDoS attack	No
NSL-KDD	2009	Canada	Generic IT	~150 MB	41	DoS, Probe, R2L, U2R	No
UNSW-NB15	2015	Australia	Generic IT	~100 MB	49	9 attack families	Partial
CIC-IDS2017	2017	Canada	Enterprise / Web	~50 GB	80+	Web, DoS, DDoS, Botnet, Brute-force	Yes (Slowloris, SlowHTTPTest)
BoT-IoT	2018	Australia	IoT	~70 GB	46	DDoS, DoS, OS scan, theft	No
CIC-IDS2018	2018	Canada	Enterprise / Web	~430 GB	80+	DDoS, Botnet, Web, Infiltration	Yes (Slowloris, SlowHTTPTest)
N-BaIoT	2018	Israel	IoT	~1.2 GB	115	Mirai, Bashlite variants	No
CICDDoS2019	2019	Canada	Cloud / Web	~24 GB	80+	13 DDoS types (SYN, UDP, DNS..)	Partial (mostly volumetric)
BENNETT-SDN (AHUJA)	2020	India	SDN testbed	~50 MB	23	TCP SYN flood, UDP flood, ICMP flood	No (all volumetric floods generated via Hping3)
LITNET-2020	2020	Lithuania	ISP / Real	~3 GB	85	12 attack categories	Partial
TON_IoT	2020	Australia	IoT / IIoT	~40 GB	44	9 categories incl. DDoS, MITM	Partial
WUSTL-EHMS	2020	USA	Healthcare	~50 MB	44	Spoofing, data injection	No
CIC-IoT2023	2023	Canada	IoT	~13 GB	47	33 attacks in 7 categories	Yes (HTTP slow)
CIC-IoMT2024	2024	Canada	Healthcare / IoMT	~6 GB	45+	18 attacks (Wi-Fi, MQTT, BLE)	No (volumetric DDoS / DoS only)
SMVDU SDN TCP-SYN	2025	India	SDN testbed	Mendeley	84	TCP-SYN flood only (single	No (single volumetric flood attack)

						attack class + benign)	
UL-ECE-DDoS-H-IoT	2025	Ireland	Healthcare-IoT (Cooja / ns-3)	~5.3 MB	Not publicly mentioned	Binary labels (0 = normal, 1 = DDoS) on two protocols: MQTT (Cooja) + UDP (ns-3)	Closer to low-rate frequency-based attack: normal nodes publish at 60-sec intervals, attack nodes at 20-sec intervals (3× anomaly, not classical Slowloris but subtler than a flood)

8. MAPPING SECTORS TO DATASETS

Fig. 7 shows the recommended mapping between sectors and datasets. Multiple datasets sometimes link to the same sector because a single dataset rarely covers all attack types of interest.

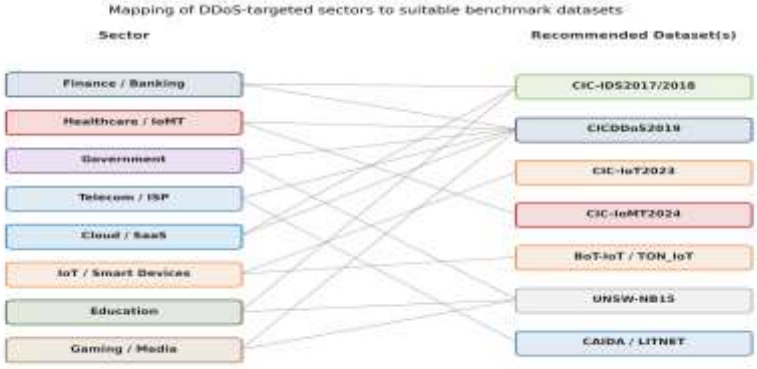

Fig 7. Recommended mapping of DDoS-targeted sectors to suitable benchmark datasets.

Table III is the practical version of this mapping, now updated to flag low-rate adequacy explicitly. Where the primary dataset lacks enough low-rate samples, the recommendation is to augment it with CIC-IDS2017 (the richest source of labelled Slowloris-family traffic in the open landscape).

TABLE III. SECTOR-DATASET MAPPING WITH LOW-RATE ADEQUACY FLAG

Sector	Primary dataset	Secondary / augment	Why this fits	Low-rate adequacy
Finance / Banking	CICDDoS2019	CIC-IDS2017 (for slow-rate)	Web-based DDoS attacks (HTTP, SYN, slow-rate) are dominant in banking; CICDDoS2019 has all major types.	Adequate only after augmenting
Healthcare / IoMT	CIC-IoMT2024	WUSTL-EHMS, CIC-IDS2017	Only IoMT-specific public dataset with Wi-Fi / MQTT / BLE attacks; WUSTL adds biometric features.	Good
Telecom / ISP	CAIDA	LITNET-2020	Both contain real (non-synthetic) backbone traffic, closer to ISP environments than testbed data.	Limited; pair with synthetic SIP traces
IoT / Smart Devices	CIC-IoT2023	BoT-IoT / TON_IoT	CIC-IoT2023 has 105 devices and 33 attacks; BoT-	Partial; augment if low-rate is target

			IoT/TON_IoT add IIoT and OS-level data.	
Energy / Utilities	TON_IoT	CIC-IoT2023	TON_IoT includes industrial IoT telemetry, the closest public proxy for SCADA-style traffic.	Limited; needs SCADA augmentation
E-commerce / Retail	CICDDoS2019	CIC-IDS2017	HTTP flood + bot-driven scraping is well-represented in CICDDoS2019 and CIC-IDS2017.	Adequate after augmenting

9. LOW-RATE DDOS COVERAGE ACROSS DATASETS

Because low-rate coverage is the key concern for our research focus, Table IV scores each dataset on three sub-criteria: the labelled low-rate attack types it contains, how realistic those samples are, and how usable the dataset is for a low-rate-only study.

TABLE IV. LOW-RATE DDOS COVERAGE SCORE PER DATASET

Dataset	Labelled low-rate attacks	Realism	Usability for low-rate study
CIC-IDS2017	Slowloris, SlowHTTPTest, Slow Read	Testbed (good)	Direct use
CIC-IDS2018	Slowloris, SlowHTTPTest	AWS-based (good)	Direct use
CIC-IoMT2024	Slow MQTT, slow HTTP	IoMT testbed (good)	Direct use (healthcare)
LITNET-2020	App-layer slow attacks	Real (excellent)	Partial - small share
UNSW-NB15	Some DoS scenarios	Testbed (ok)	Needs augmentation
TON_IoT	App-layer DoS	Testbed (ok)	Needs augmentation
CIC-IoT2023	Few HTTP slow scenarios	Testbed (ok)	Needs augmentation
CICDDoS2019	Mostly volumetric reflection	Testbed (good)	Augment heavily
BoT-IoT	None labelled	Testbed	Not suitable
N-BaIoT	None labelled	Testbed	Not suitable
NSL-KDD	None labelled	Outdated	Not suitable
CAIDA	None labelled	Real (good)	Not suitable
WUSTL-EHMS	None labelled (focus: spoofing)	Testbed (ok)	Not suitable for DDoS work

Three datasets stand out for low-rate DDoS research: CIC-IDS2017, CIC-IDS2018 and CIC-IoMT2024. CIC-IDS2017 is the gold standard for slow-rate web attacks; CIC-IDS2018 extends the same idea to a cloud (AWS) environment; and CIC-IoMT2024 is the most relevant choice when the research target is healthcare and IoMT. A pragmatic strategy used by several recent papers [9], [25], [30] is to combine CICDDoS2019 (for volumetric variety) with CIC-IDS2017 (for low-rate variety) and train one model on the merged set, ensuring that the final classifier sees both attack regimes.

10. THE INDIAN PERSPECTIVE

Until now this survey has stayed mostly with the global picture. But this work is being carried out in India, and the picture in India is not quite the same as the global one. This section pulls together what the recent data tells us about DDoS in the Indian context.

10.1 India on the global DDoS map

India has moved up the list of most-attacked countries very quickly. In the first half of 2025 India received about 12.6 percent of all DDoS attacks reported worldwide, up from 9.2 percent a year earlier [38]. That works out to roughly a 37 percent rise in twelve months, and it puts India second only to the United States. China, which used to be ahead of India, is now third. The financial sector is bearing the heaviest pressure inside the country. From January to June 2025, banks and other financial institutions in India faced

an average of 4.1 million attacks every month a 172 percent jump in DDoS attacks during peak banking operations compared to the previous year [39]. The Reserve Bank of India has confirmed 248 data breaches at scheduled commercial banks over a four-year window. Most security analysts believe the real count is higher. Many incidents never get reported in the first place.

10.2 Recent attacks on Indian healthcare

The last few years have been rough for Indian hospitals. The one everyone remembers is the AIIMS Delhi ransomware attack on November 23, 2022. The e-hospital server went down. The institution had to fall back to manual patient records, hand-written notes and paper billing for several weeks [40]. News reports estimate that the medical records of around 40 million patients may have been exposed [41]. Of AIIMS's hundred physical servers, five were successfully breached by the attackers [42]. AIIMS Delhi was not the only target. A few days later, Safdarjung Hospital in Delhi reported a similar cyber incident, and around the same time the records of about 150,000 patients from Sree Saran Medical Centre in Tirupur, Tamil Nadu, appeared on the dark web [43]. AIIMS was attacked a second time in 2023, this time with a more sophisticated malware campaign that was described as targeted [44]. None of these were classical low-rate DDoS attacks. But they show how exposed Indian healthcare infrastructure is, a successful low-rate attack on a hospital portal or on an IoMT device would be even harder to notice than the ransomware incidents that have already happened.

10.3 What public DDoS datasets exist in India

Honestly not many, and none that match a healthcare low-rate research focus. The Bennett-SDN dataset (Ahuja, Singal and Mukhopadhyay, Bennett University, 2020) is the most cited Indian DDoS dataset by a wide margin. It is built using the Mininet emulator with a Ryu SDN controller, includes ten different topologies, and contains 23 flow-level features per record [35]. The attacks covered are TCP SYN flood, UDP flood and ICMP flood all volumetric. The dataset is freely available on Mendeley. The SMVDU SDN TCP-SYN dataset (Kumar and Gupta, Shri Mata Vaishno Devi University, 2025) is a much smaller dataset published recently in Data in Brief. It targets only the TCP-SYN flood attack in an SDN testbed using Mininet and hping3 [36]. Useful as a single-attack benchmark but limited in scope. Put together they are volumetric in their attack types. India is the second-most attacked country in the world. Indian healthcare is being hit, repeatedly, by sophisticated attackers. Indian regulators want detection research and have made incident reporting mandatory within six hours. But there is no public Indian dataset that combines healthcare IoMT traffic with labelled low-rate DDoS samples.

11. FEASIBILITY AND IMPLEMENTATION NOTES

11.1 Size and storage

CIC-IDS2018 at ~430 GB is heavy for a low configuration laptop. The CSV-flow versions of CIC-IDS2017 (~50 GB) and CIC-IoMT2024 (~6 GB) are more realistic starting points.

11.2 Class imbalance for low-rate samples

Low-rate samples are a small minority class in every dataset that contains them. In CIC-IDS2017, for instance, Slowloris is a tiny fraction of total flows. SMOTE-style oversampling or focal-loss training is usually necessary, and reporting F1 and false positive rate (not just accuracy) is essential [7].

11.3 Synthetic versus real traffic

Most low-rate samples come from testbeds, which means they may not capture the full noise of real traffic. Where possible, validating on at least one real dataset such as LITNET-2020 or by collecting a small real trace from a cooperating environment increases credibility [13], [21].

12. CONCLUSION

Low-rate DDoS is the quieter half of the DDoS problem, and the choice of dataset matters even more for low-rate research than it does for volumetric work. In this paper we brought sector analysis and dataset selection together, with low-rate DDoS as the central concern. Ten sectors were reviewed with their specific low-rate variants. Various datasets were reviewed, with a dedicated table scoring each one on low-rate coverage. The mapping shows that, for healthcare-focused low-rate DDoS research, CIC-IoMT2024 supplemented by CIC-IDS2017 is currently the best practical starting point. For broader web or cloud research the combination of CICDDoS2019 (for volumetric diversity) and CIC-IDS2017 (for low-rate diversity) is the most balanced option.

13. ACKNOWLEDGEMENT

This research was supported by a grant from ANRF-PAIR. We are extremely grateful for their financial backing, which allowed us to do this research work. We also wish to acknowledge the support provided by IISC Team of ANRF-PAIR and the editorial suggestions offered by our anonymous reviewers.

14. REFERENCES

1. Cloudflare, "Q4 2025 DDoS Threat Report," Cloudflare blog, Jan. 2026. [Online]. Available: https://blog.cloudflare.com/ddos-threat-report-for-2025-q4/
2. StormWall, "DDoS Attacks Report 2024," StormWall, Jan. 2025. [Online]. Available: https://stormwall.network/resources/blog/ddos-report-2024
3. NETSCOUT, "DDoS Threat Intelligence Report 2H 2024," NETSCOUT, 2025. [Online]. Available: https://www.netscout.com/threatreport
4. Akamai, "State of the Internet / Security: Defenders' Guide 2024," Akamai Technologies, 2024. [Online]. Available: https://www.akamai.com/security-research/the-state-of-the-internet

5. MazeBolt, "DDoS Attack Statistics and Trends," MazeBolt Technologies, 2024. [Online]. Available: https://mazebolt.com/blog/ddos-attack-statistics/
6. S. T. Zargar, J. Joshi and D. Tipper, "A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks," IEEE Commun. Surv. Tutor., vol. 15, no. 4, pp. 2046–2069, 2013.
7. A. Bhardwaj, V. Mangat and R. Vig, "Hyperband Tuned Deep Neural Network with Well Posed Stacked Sparse AutoEncoder for Detection of DDoS Attacks in Cloud," IEEE Access, vol. 8, pp. 181916–181929, 2020.
8. M. Hassan et al., "A Systematic Review of Cyber Threats in Healthcare," J. Med. Syst., vol. 47, no. 8, 2023.
9. A. Al-Shenbary et al., "A Comprehensive Survey of Intrusion Detection Systems," Appl. Comput. Intell. Soft Comput., vol. 2026, art. no. 9123456, 2026.
10. Y. Mirsky et al., "Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection," NDSS, 2018.
11. I. Sharafaldin, A. H. Lashkari and A. A. Ghorbani, "Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization," ICISSP, pp. 108–116, 2018.
12. I. Sharafaldin, A. H. Lashkari, S. Hakak and A. A. Ghorbani, "Developing Realistic Distributed Denial of Service (DDoS) Attack Dataset and Taxonomy," ICCST, pp. 1–8, 2019. [Online]. Available: https://www.unb.ca/cic/datasets/ddos-2019.html
13. CAIDA, "DDoS Attack 2007 Dataset," 2007. [Online]. Available: https://www.caida.org/catalog/datasets/ddos-20070804_dataset/
14. N. Moustafa and J. Slay, "UNSW-NB15: A Comprehensive Data Set for Network Intrusion Detection Systems," MilCIS, pp. 1–6, 2015. [Online]. Available: https://research.unsw.edu.au/projects/unsw-nb15-dataset
15. CIC, "CICIDS2017 / CSE-CIC-IDS2018 Datasets," University of New Brunswick, 2018. [Online]. Available: https://www.unb.ca/cic/datasets/ids-2018.html
16. N. Koroniotis, N. Moustafa, E. Sitnikova and B. Turnbull, "Towards the Development of Realistic Botnet Dataset in the Internet of Things for Network Forensic Analytics: Bot-IoT Dataset," Future Gener. Comput. Syst., vol. 100, pp. 779–796, 2019.
17. N. Moustafa, "A New Distributed Architecture for Evaluating AI-Based Security Systems at the Edge: Network TON_IoT Datasets," Sustain. Cities Soc., vol. 72, art. no. 102994, 2021. [Online]. Available: https://research.unsw.edu.au/projects/toniot-datasets
18. S. Dadkhah et al., "CICIoT2023: A Real-Time Dataset and Benchmark for Large-Scale Attacks in IoT Environment," Sensors, vol. 23, no. 13, 2023. [Online]. Available: https://www.unb.ca/cic/datasets/iotdataset-2023.html
19. S. Dadkhah et al., "CIC IoMT 2024: A Benchmark Dataset for IoMT Security Research," IEEE Dataport, 2024. [Online]. Available: https://www.unb.ca/cic/datasets/iomt-dataset-2024.html
20. A. Hady, A. Ghubaish, T. Salman, D. Unal and R. Jain, "Intrusion Detection System for Healthcare Systems Using Medical and Network Data: A Comparison Study," IEEE Access, vol. 8, pp. 106576–106584, 2020. [Online]. Available: https://www.cse.wustl.edu/~jain/ehms/
21. R. Damasevicius et al., "LITNET-2020: An Annotated Real-World Network Flow Dataset for Network Intrusion Detection," Electronics, vol. 9, no. 5, art. no. 800, 2020.
22. H. Sun, J. C. S. Lui and D. K. Y. Yau, "Defending Against Low-Rate TCP Attacks: Dynamic Detection and Protection," IEEE ICNP, pp. 196–205, 2004.
23. A. Kuzmanovic and E. W. Knightly, "Low-Rate TCP-Targeted Denial of Service Attacks: The Shrew vs. the Mice and Elephants," ACM SIGCOMM, pp. 75–86, 2003.
24. R. Hansen (RSnake), "Slowloris HTTP DoS," Tech. Rep., 2009. [Online]. Available: https://github.com/gkbrk/slowloris
25. S. Mohammed et al., "A Comprehensive Review on IoMT Security: Challenges and Countermeasures," J. King Saud Univ. – Comput. Inf. Sci., vol. 38, no. 2, 2026.
26. K. Shamsolmoali and M. Zareapoor, "Statistical-Based Filtering System Against DDoS Attacks in Cloud Computing," Int. Conf. Adv. Comput., pp. 1234–1239, 2014.
27. J. Tang et al., "Low-Rate DoS Attacks against SIP Servers: Detection and Mitigation," Comput. Secur., vol. 78, pp. 23–37, 2018.
28. T. Mahjabin, Y. Xiao, G. Sun and W. Jiang, "A Survey of Distributed Denial-of-Service Attack, Prevention, and Mitigation Techniques," Int. J. Distrib. Sens. Netw., vol. 13, no. 12, 2017.
29. M. Antonakakis et al., "Understanding the Mirai Botnet," USENIX Security Symposium, pp. 1093–1110, 2017.
30. Y. Tang, K. Mei and J. Chen, "A Survey of Low-Rate DDoS Detection Techniques," IEEE Access, vol. 11, pp. 23981–24001, 2023.
31. J. Mirkovic and P. Reiher, "A Taxonomy of DDoS Attack and DDoS Defense Mechanisms," ACM SIGCOMM Comput. Commun. Rev., vol. 34, no. 2, pp. 39–53, 2004.
32. R. K. C. Chang, "Defending against Flooding-Based Distributed Denial-of-Service Attacks: A Tutorial," IEEE Commun. Mag., vol. 40, no. 10, pp. 42–51, 2002.
33. E. Cambiaso, G. Papaleo, and M. Aiello, "Slow DoS attacks: Definition and categorisation," International Journal of Trust Management in Computing and Communications, vol. 1, no. 3–4, pp. 300–319, 2013.
34. S. S. Bamber, A. V. R. Katkuri, S. Sharma, and M. Angurala, "A hybrid CNN-LSTM approach for intelligent cyber intrusion detection system," Computers & Security, vol. 148, art. 104146, 2025. doi: 10.1016/j.cose.2024.104146.
35. N. Ahuja, G. Singal and D. Mukhopadhyay, "DDOS Attack SDN Dataset," Mendeley Data, V1, 2020. DOI: 10.17632/jxpfjc64kr.1. [Online]. Available: https://data.mendeley.com/datasets/jxpfjc64kr/1
36. S. Kumar and S. Gupta, "SDN TCP-SYN Dataset: A Dataset for TCP-SYN Flood DDoS Attack Detection in Software-Defined Networks," Data in Brief, vol. 58, art. no. 111314, Jan. 2025. DOI: 10.1016/j.dib.2025.111314. [Online]. Available: https://data.mendeley.com/datasets/236bd4cjmk/2

37. M. Akhi, C. Eising and L. L. Dhirani, "Datasets for Distributed Denial-of-Service Detection in Healthcare Internet of Things Environments," Data in Brief, vol. 56, art. no. 112222, Nov. 2025. DOI: 10.1016/j.dib.2025.112222. [Online]. Available: https://data.mendeley.com/datasets/2bw34ght8b/2
38. StormWall, "DDoS in H1 2025: Analytics and Statistics," StormWall blog, Dec. 2025. [Online]. Available: https://stormwall.network/resources/blog/ddos-attack-statistics-h1-2025
39. P. Mali, "Data Breaches in India's Banking Sector in 2025: A Comprehensive Analysis," Cyber Law Consulting,Nov.2025.[Online].Available:https://www.cyberlawconsulting.com/Data_Breaches_in_India_Banking_Sector_in_2025_A_Comprehensive_Analysis.php
40. Healthcare IT News, "AIIMS Delhi Turns Manual Following Ransomware Attack," Healthcare IT News Asia, Dec. 2022. [Online]. Available: https://www.healthcareitnews.com/news/asia/aiims-delhi-turns-manual-following-ransomware-attack
41. M. Singh, "India's AIIMS Hit by Outages After Cyberattack," TechCrunch, Nov. 24, 2022. [Online]. Available: https://techcrunch.com/2022/11/24/aiims-india-cyber-attack/
42. NLIU Cell for Studies in Intellectual Property Rights, "The AIIMS Cyber-Attack and India's Dilapidated Cyber-security Infrastructure," NLIU CSIPR, 2023. [Online]. Available: https://csipr.nliu.ac.in/news-updates/the-aiims-cyber-attack-and-indias-dilapidated-cyber-security-infrastructure/
43. Healthcare IT News, "Not Just AIIMS Delhi: Safdarjung Hospital Also Reports Cyberattack," Healthcare IT News Asia, Dec. 2022. [Online]. Available: https://www.healthcareitnews.com/news/asia/not-just-aiims-delhi-safdarjung-hospital-also-reports-cyberattack
44. CyberPeace Foundation, "Cyber Attack Alert! AIIMS Attacked Again," CyberPeace blog, Aug. 2023. [Online]. Available: https://cyberpeace.org/resources/blogs/cyber-attack-alert-aiims-attacked-again