CSCI E-11

Class 2

February 3, 2026

Risks and Risk Management

Harvard
Extension School
HARVARD DIVISION OF
CONTINUING EDUCATION

Administrivia

- Quizzes
  - Some questions have multiple correct answers!
  - We've set things to credit you with your highest score, not simply your latest. (Labs work this way too.)
  - Answers may shuffle between tries.
- Getting Ahead
  - All labs and readings are available (but the lectures will provide context).
  - The next few reading quizzes are available.
- Perusall
  - We've experimented with splitting the class up into groups. Let us know what you think.
- Participation

One of the best-known client/server protocols is the Hypertext Transfer Protocol.

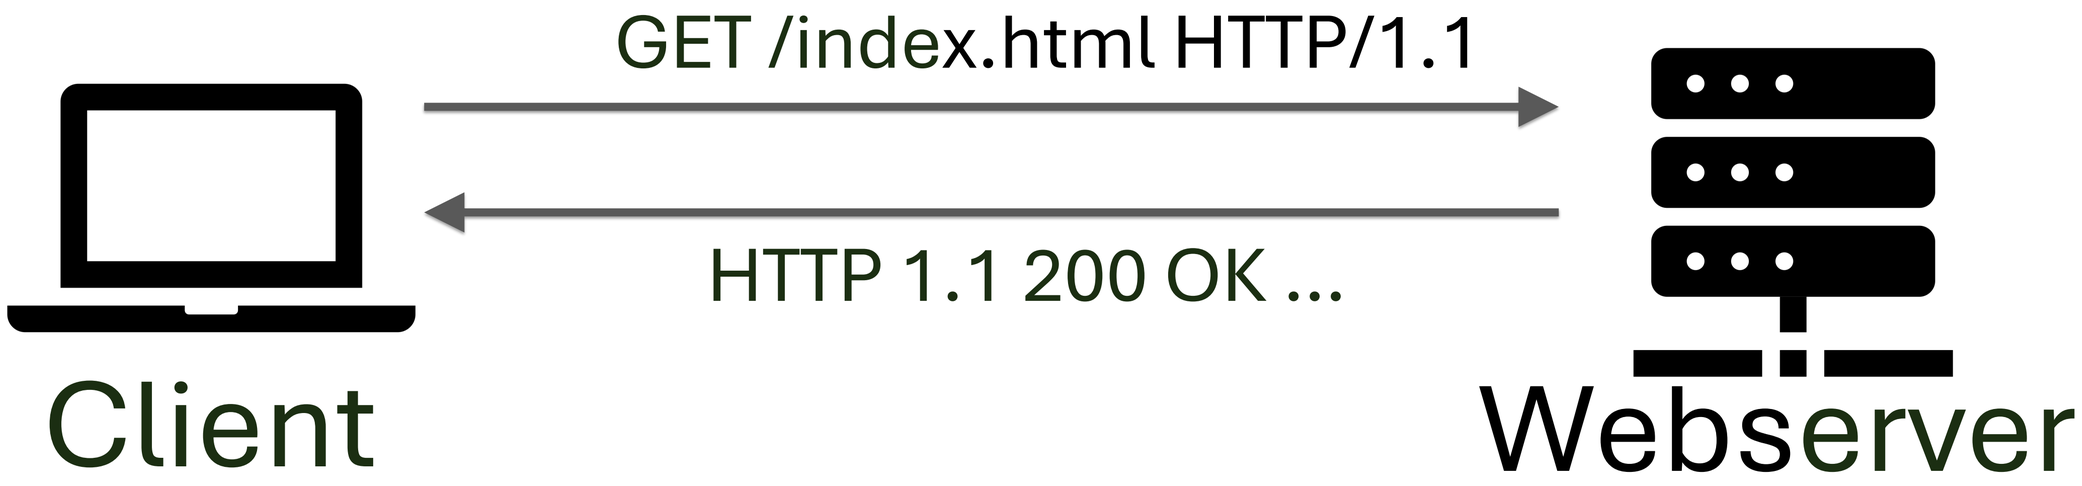

HTTP is sent in the clear, so is vulnerable to interception and surveillance.

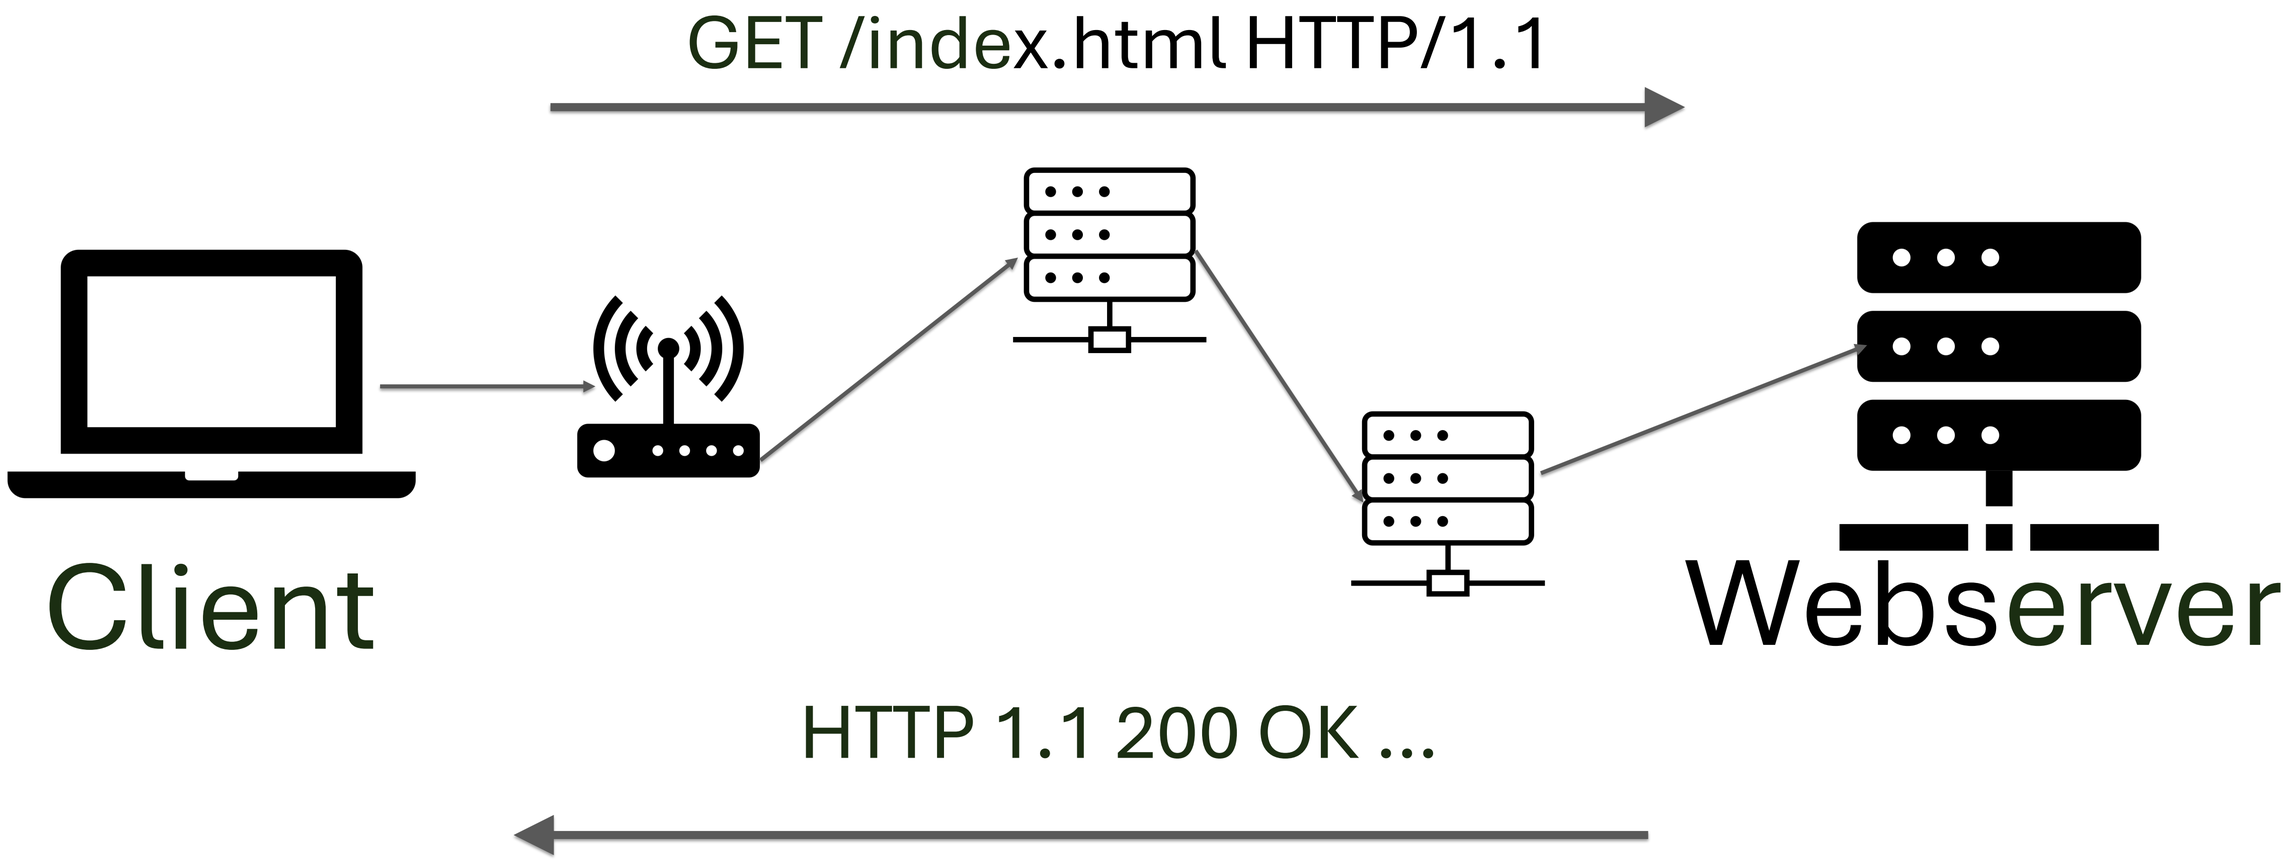

“Certificates” enable encryption and authentication.

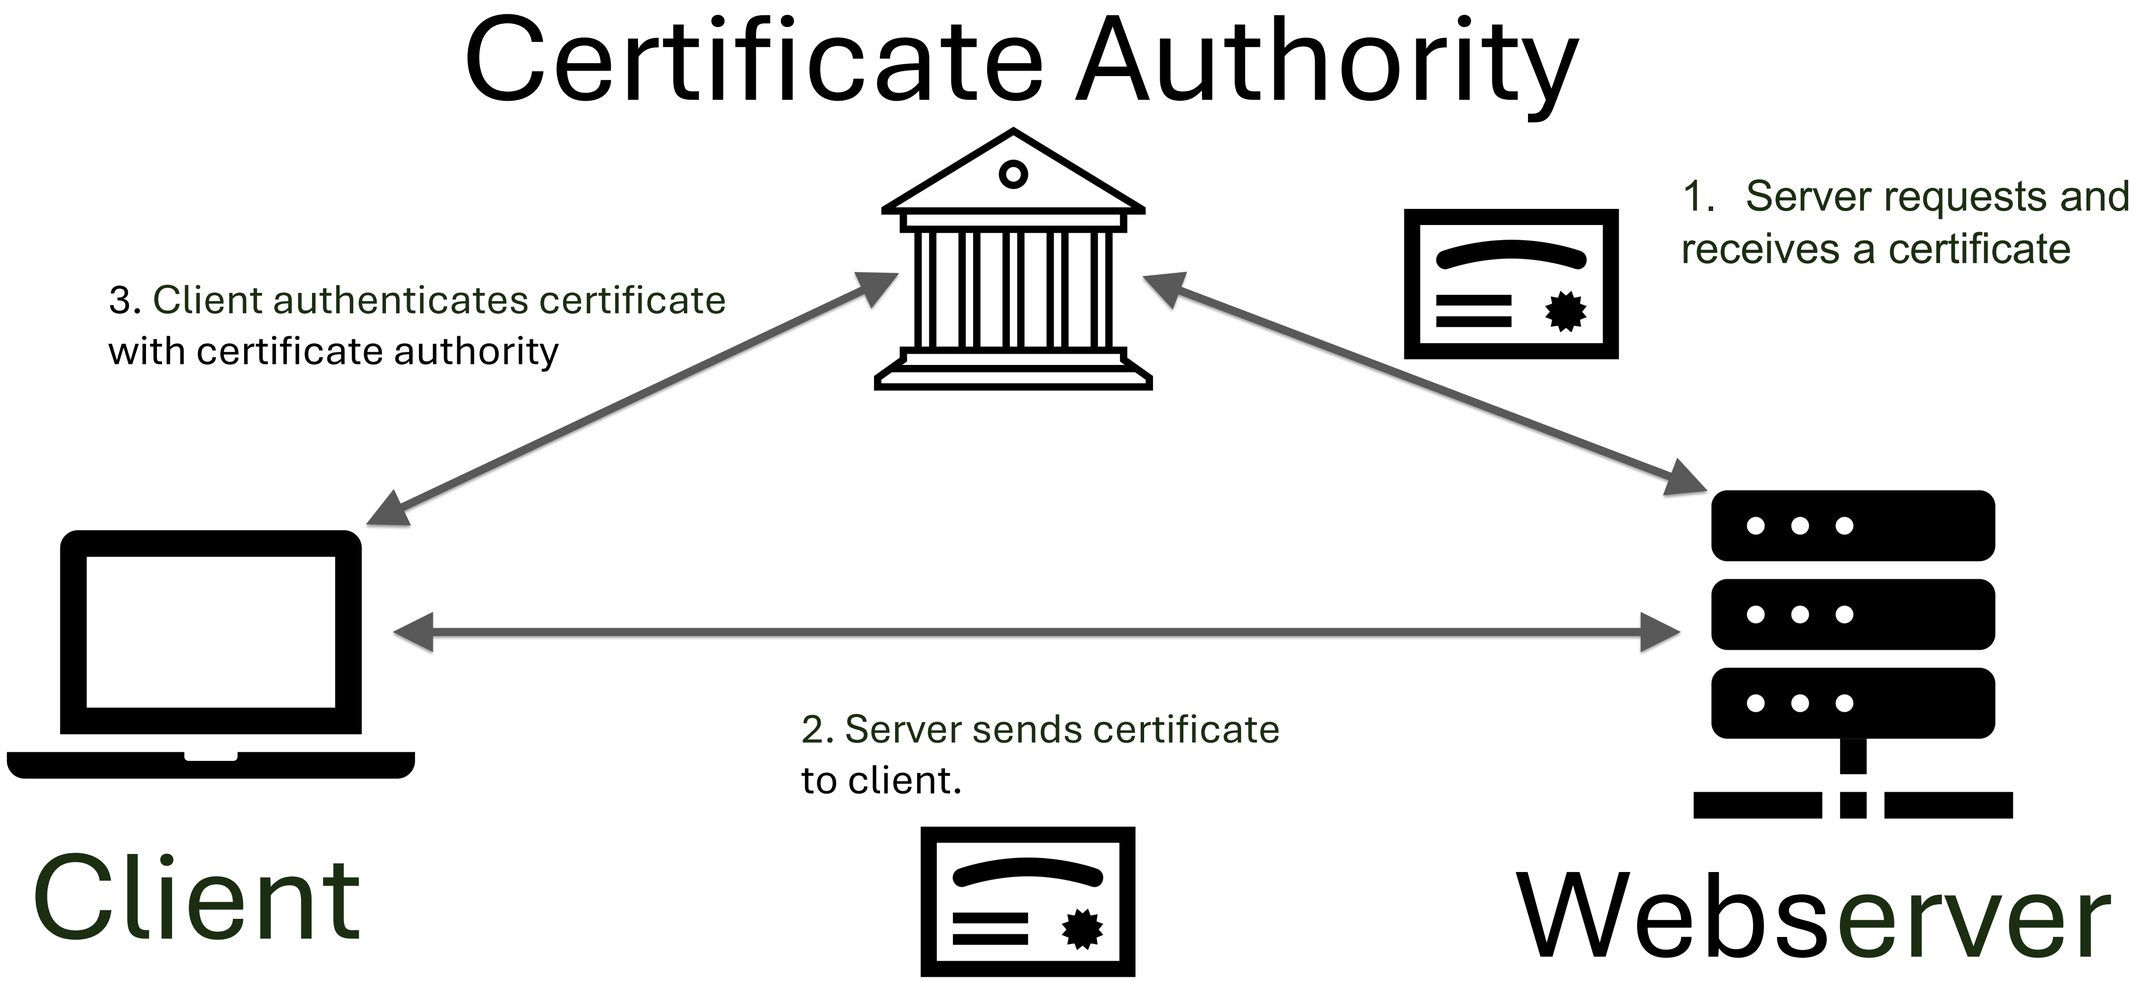

Your browser will tell you if you're using a valid certificate.

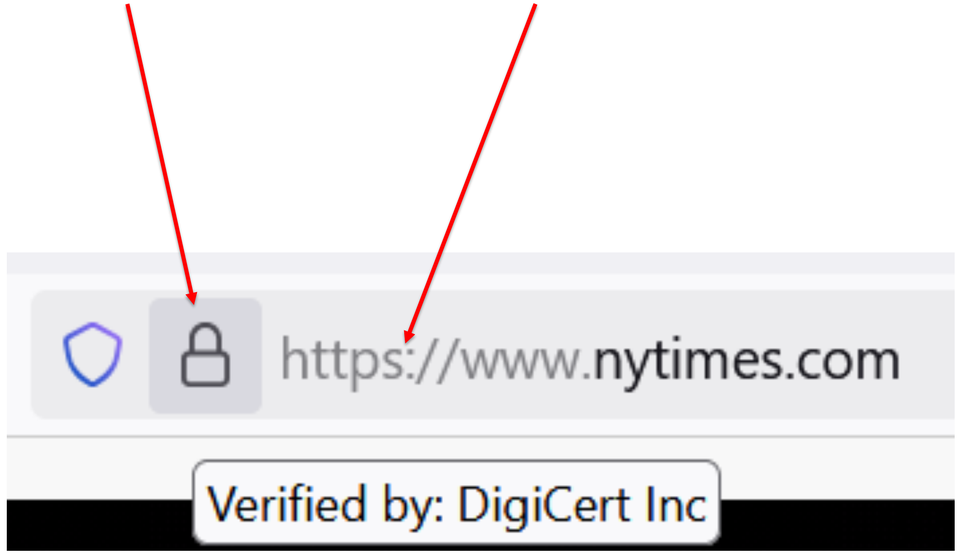

Set your browser to use only HTTPS connections for all websites.

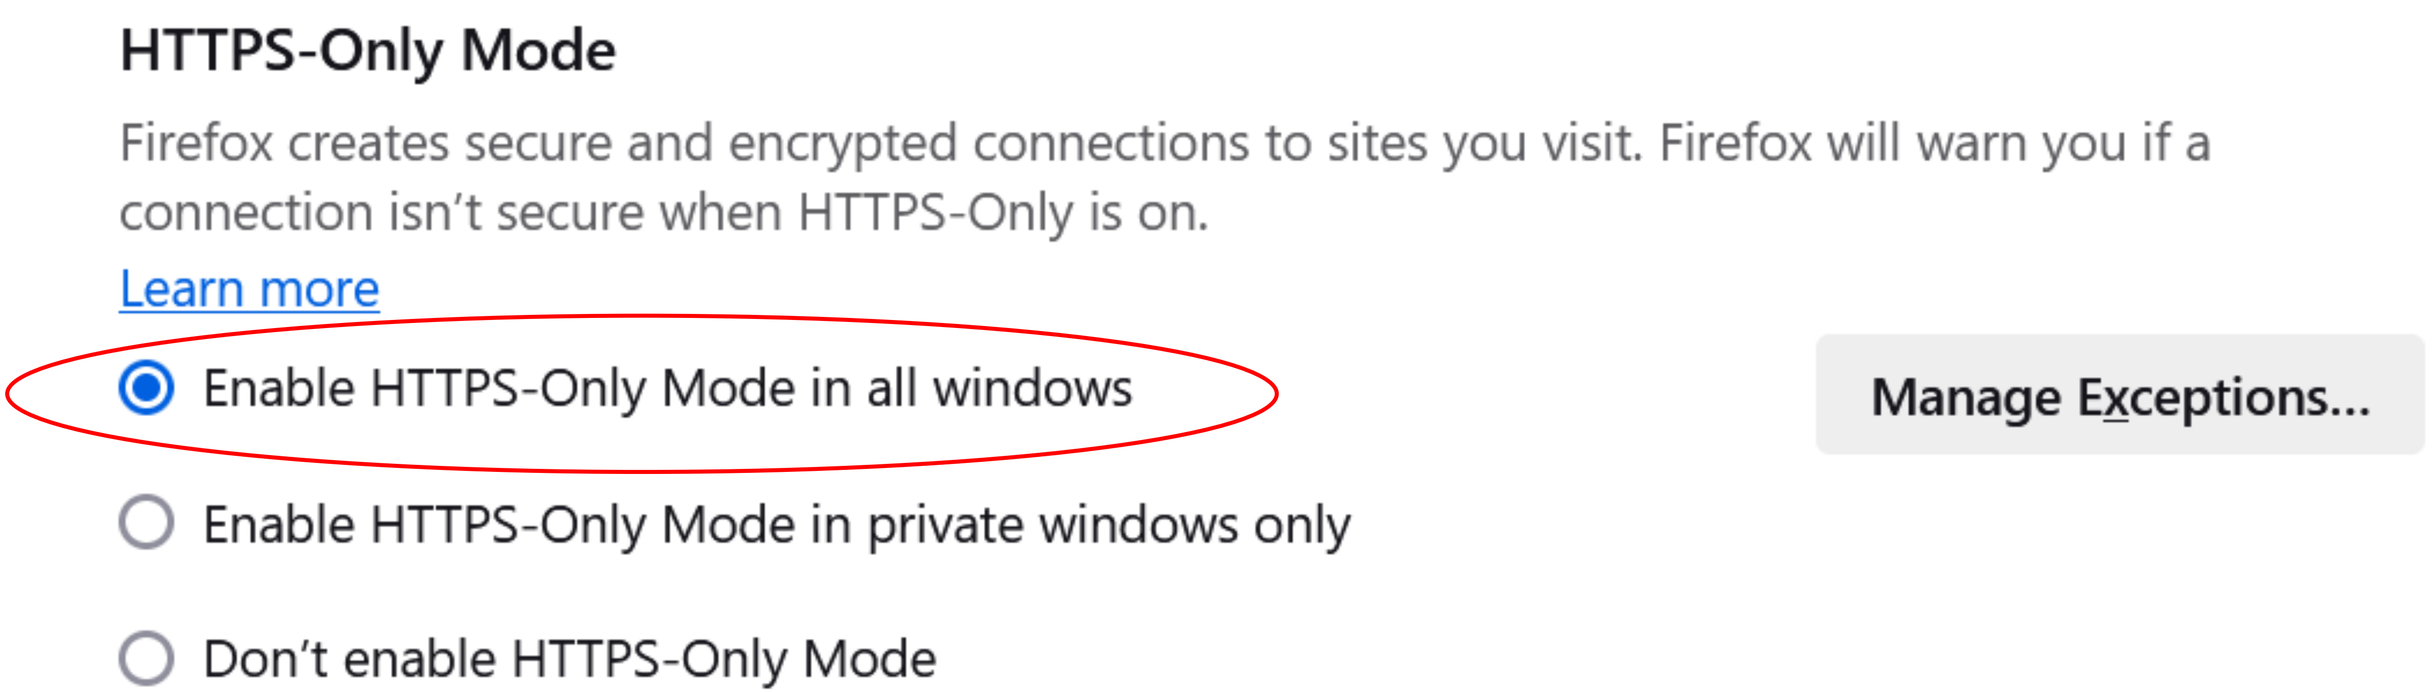

The Certificate Authority issues the Certificate to the legitimate domain.

1. Server requests and receives a certificate

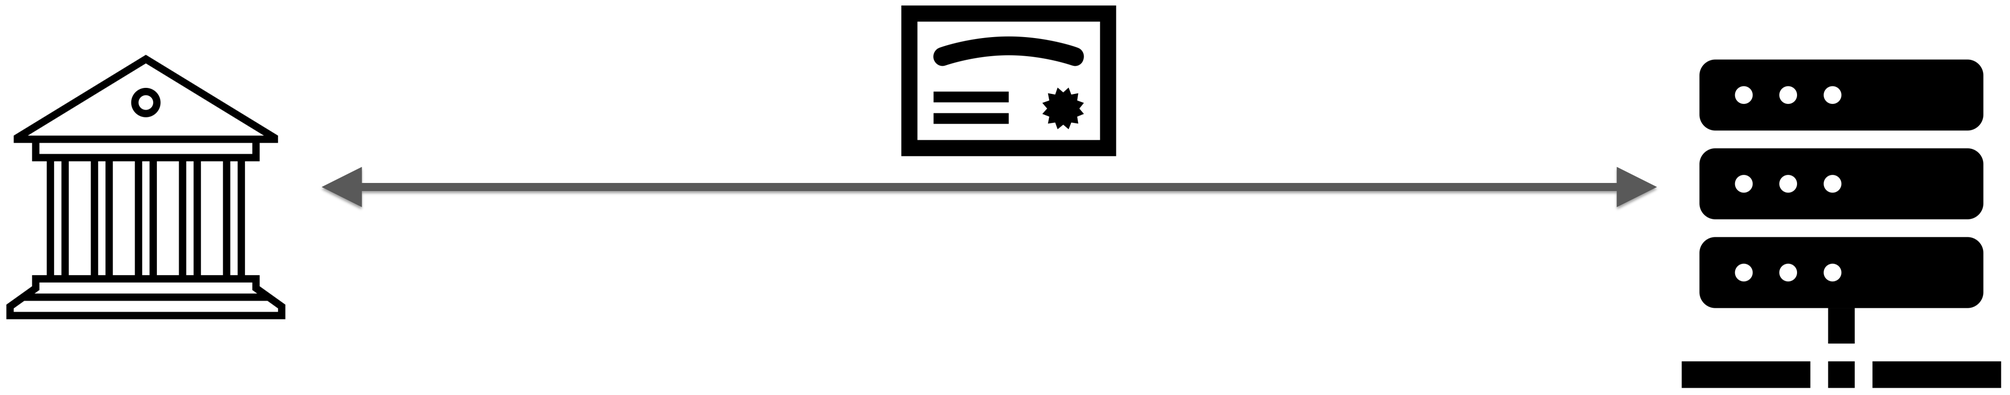

Certificate Authority

Webserver

The Certificate Authority issues the Certificate for the domain.

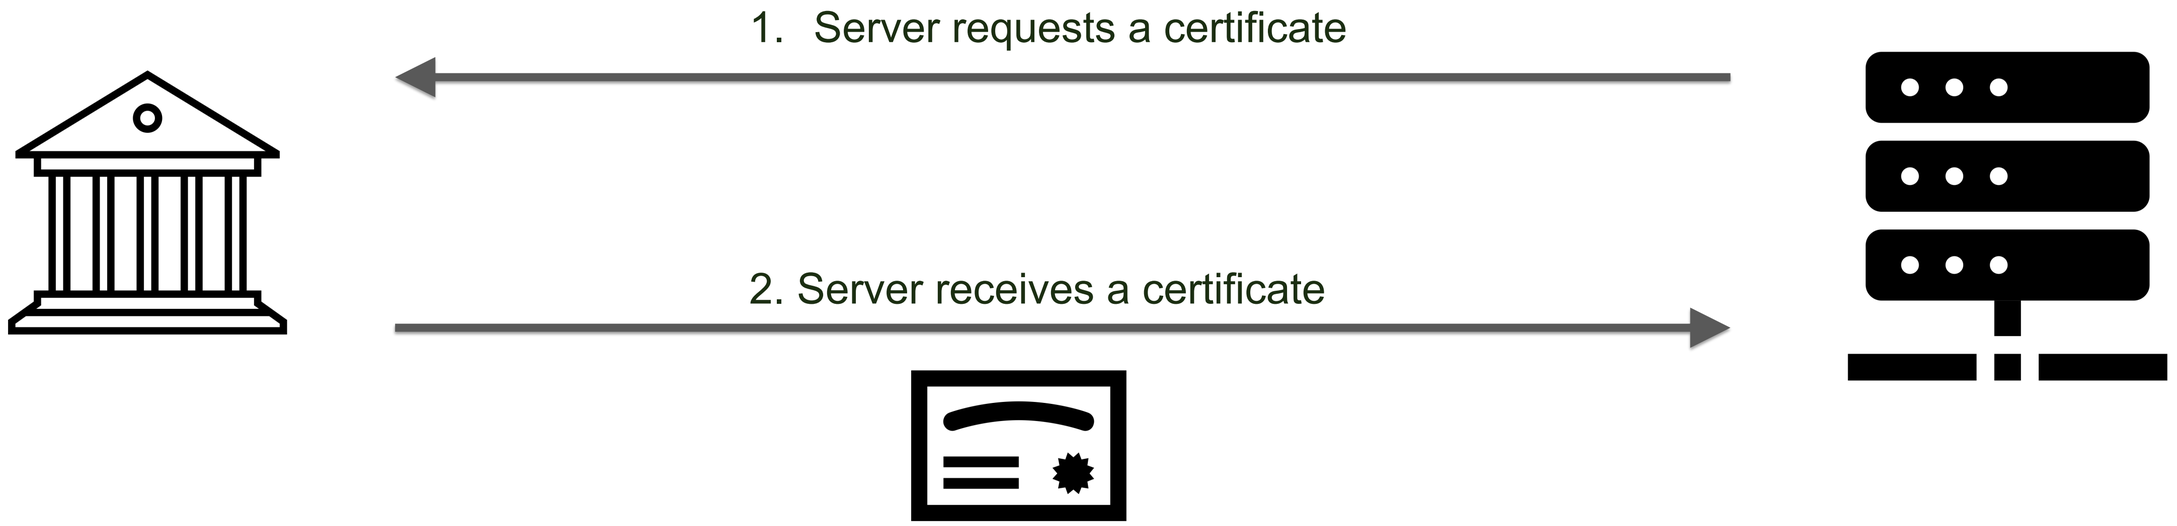

ACME makes sure that Certificates only issue to the legitimate domain.

Server requests a token

CA transmits a token

GET http://<DOMAIN>/.well-known/acme-challenge/<TOKEN>

Server returns token

Server receives a certificate

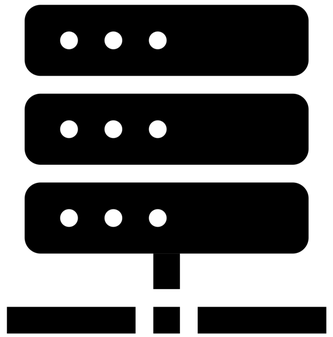

Certificate Authority

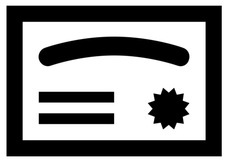

Webserver

Cloudflare is an Internet security company that provides proxies that protect the origin server. This product is a “Web Application Firewall.”

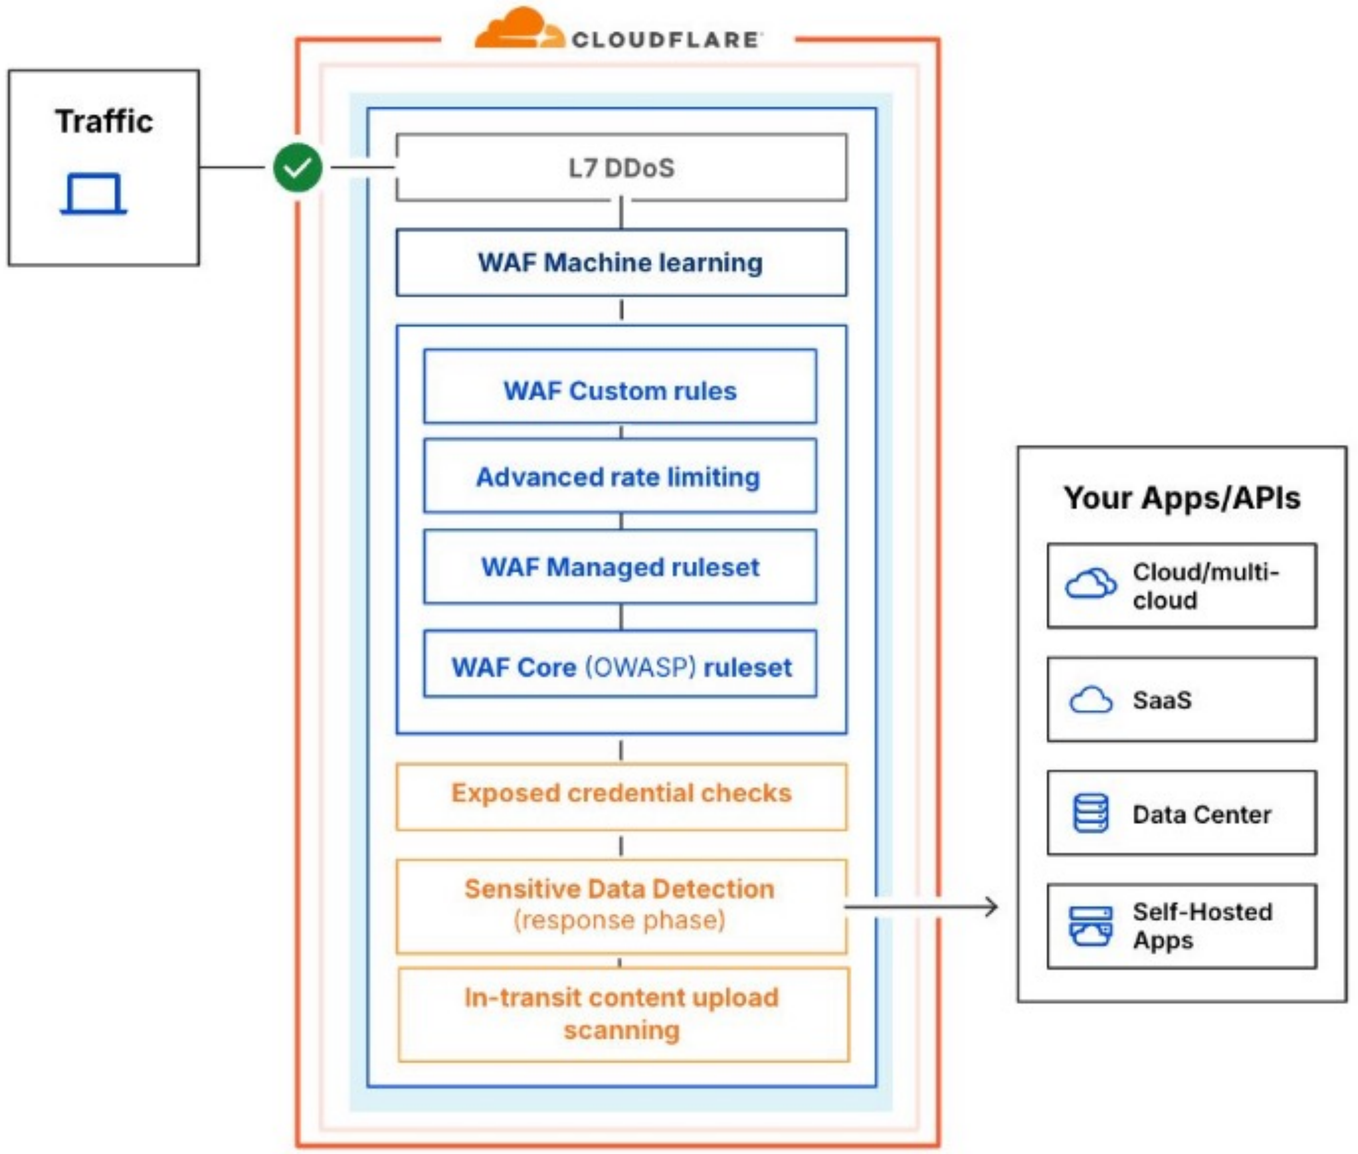

An exploit allowed malicious requests of tokens to get through the firewall.

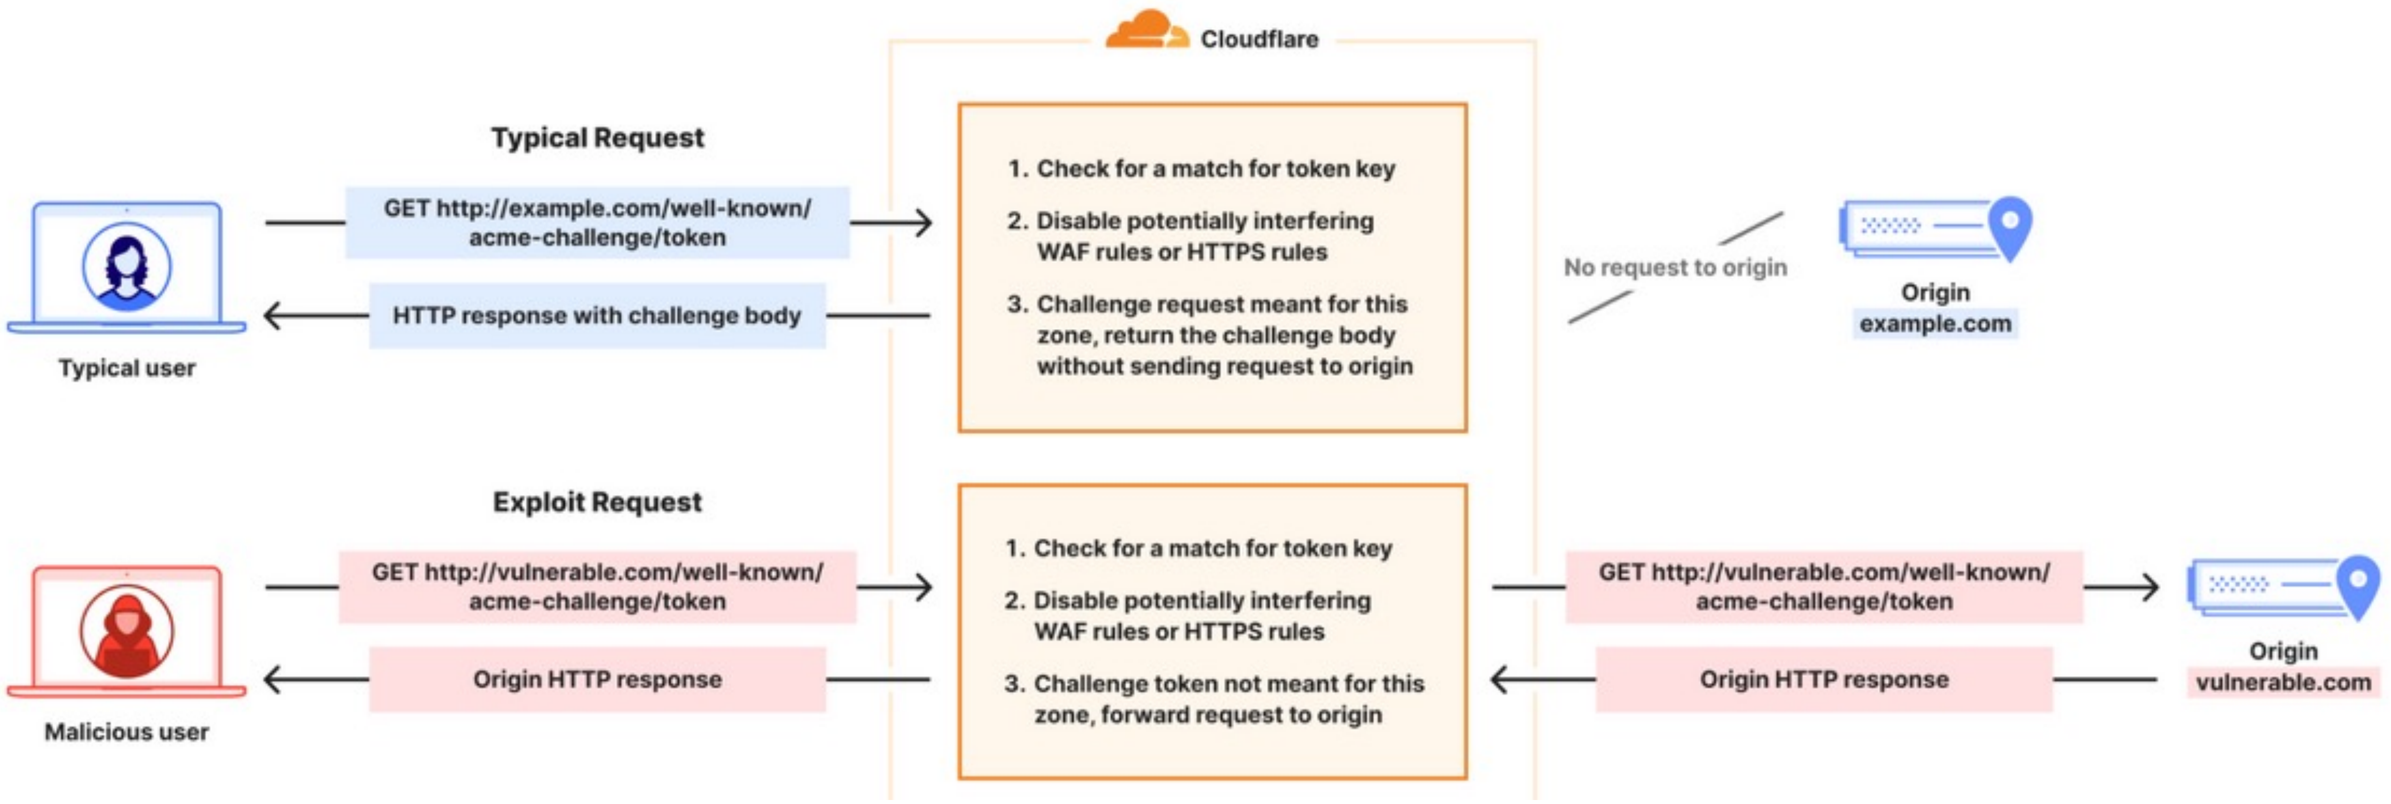

FearsOff, a security research firm, identified the vulnerability.

Under a normal posture, actuator endpoints sit behind the WAF and internal network controls. Reaching the application through the ACME path changed that boundary. Using a well known traversal quirk in some servlet stacks (`..;/`), the request could land on `/actuator/env` and return process environment and configuration. That data often includes sensitive values - database URLs, API tokens, cloud keys - and it materially raises the blast radius of any mistake in the origin.

https://fearsoff.org/research/cloudflare-acme

The exploit was submitted to Cloudflare's Bug Bounty program.

Program highlights

Open Scope	Rewards reports for all owned assets based on impact, even if not listed in scope.
Gold Standard Safe Harbor	Adheres to Gold Standard Safe Harbor.
Platform Standards	Fully compliant with Platform Standards.
Top Response Efficiency	This program's response efficiency is above 90%.

Managed by HackerOne | Collaboration Enabled | Includes Retesting

Rewards summary

Last updated on October 4, 2024. View changes

Each severity lists the 90-day average bounty and the percentage of total resolved reports, if applicable.

Low	Medium	High	Critical
Avg. bounty $300 28.70% submissions	Avg. bounty $570 37.16% submissions	Avg. bounty $1,582 24.77% submissions	Avg. bounty $3,900 9.37% submissions
$250–$500	$500–$750	$1,000–$3,000	$4,000–$10,000

This week's reading: Peter G. Neumann

viewpoints

DOI:10.1145/3532631

Peter G. Neumann

Inside Risks

Toward Total-System Trustworthiness

Considering how to achieve the long-term goal to systemically reduce risks.

COMMUNICATIONS' INSIDE RISKS columns have long stressed the importance of total-system awareness of riskful situations, some of which may be very difficult to identify in advance. Specifically, the desired properties of the total system should be specified as requirements. Those desired properties are called emergent properties, because they often cannot be derived solely from lower-layer component properties, and appear only with respect to the total system. Unfortunately, additional behavior of the total system may arise—which either defeats the ability to satisfy the desired properties, or demonstrates that the set of required properties was improperly specified.

In this column, I consider some cases in which total-system analysis is of vital importance, but generally very difficult to achieve with adequate assurance. Relevant failures may result from one event or even a combination of problems in hardware, software, networks, operational environments, and of course actions by administrators, users, and misusers. All of the interactions among these entities need to be considered, evaluated, and if potentially deleterious, controlled by whatever means available. The problem to be confronted here is trying to analyze an entire system as a composition of its components, rather than just considering its components individually. In many cases, system failures tend to arise within the interactions and interdependencies among these components, depending on whether a system was designed modularly to minimize disruptive dependencies, with each module carefully specified.

Addressing this problem is a daunting endeavor, even for seasoned developers of critical systems. Whether explicitly defined or implicit, total-system requirements may be highly interdisciplinary, including stringent life-critical requirements for human safety; system survivability with reliability, robustness, resilience, recovery, and fault tolerance; many aspects of security and integrity; guaranteed real-time performance; forensics-worthy accountability, high-integrity evidence, and sound real-time and retrospective analysis; defenses against a wide ranges of physical, electronic, and other adversities; and coverage of numerous potential risks. Ideally, requirements should be very carefully specified at various architectural layers, preferably formally as much as possible in newly developed systems, and especially in particularly vulnerable components. Although this is usually not applicable to legacy systems, it is stated here as a farsighted goal for future developments.

IMAGE FROM SHUTTERSTOCK.COM

Discussion themes:

- Emergent properties
- AI as black box
- Systems change and old guarantees don't hold
- Legacy
- Trade-offs
- Catastrophic vs. graceful failure
- "Long-term"
- Over-reliance (e.g. on encryption)
- Is it even possible?

Welcome Survey Results

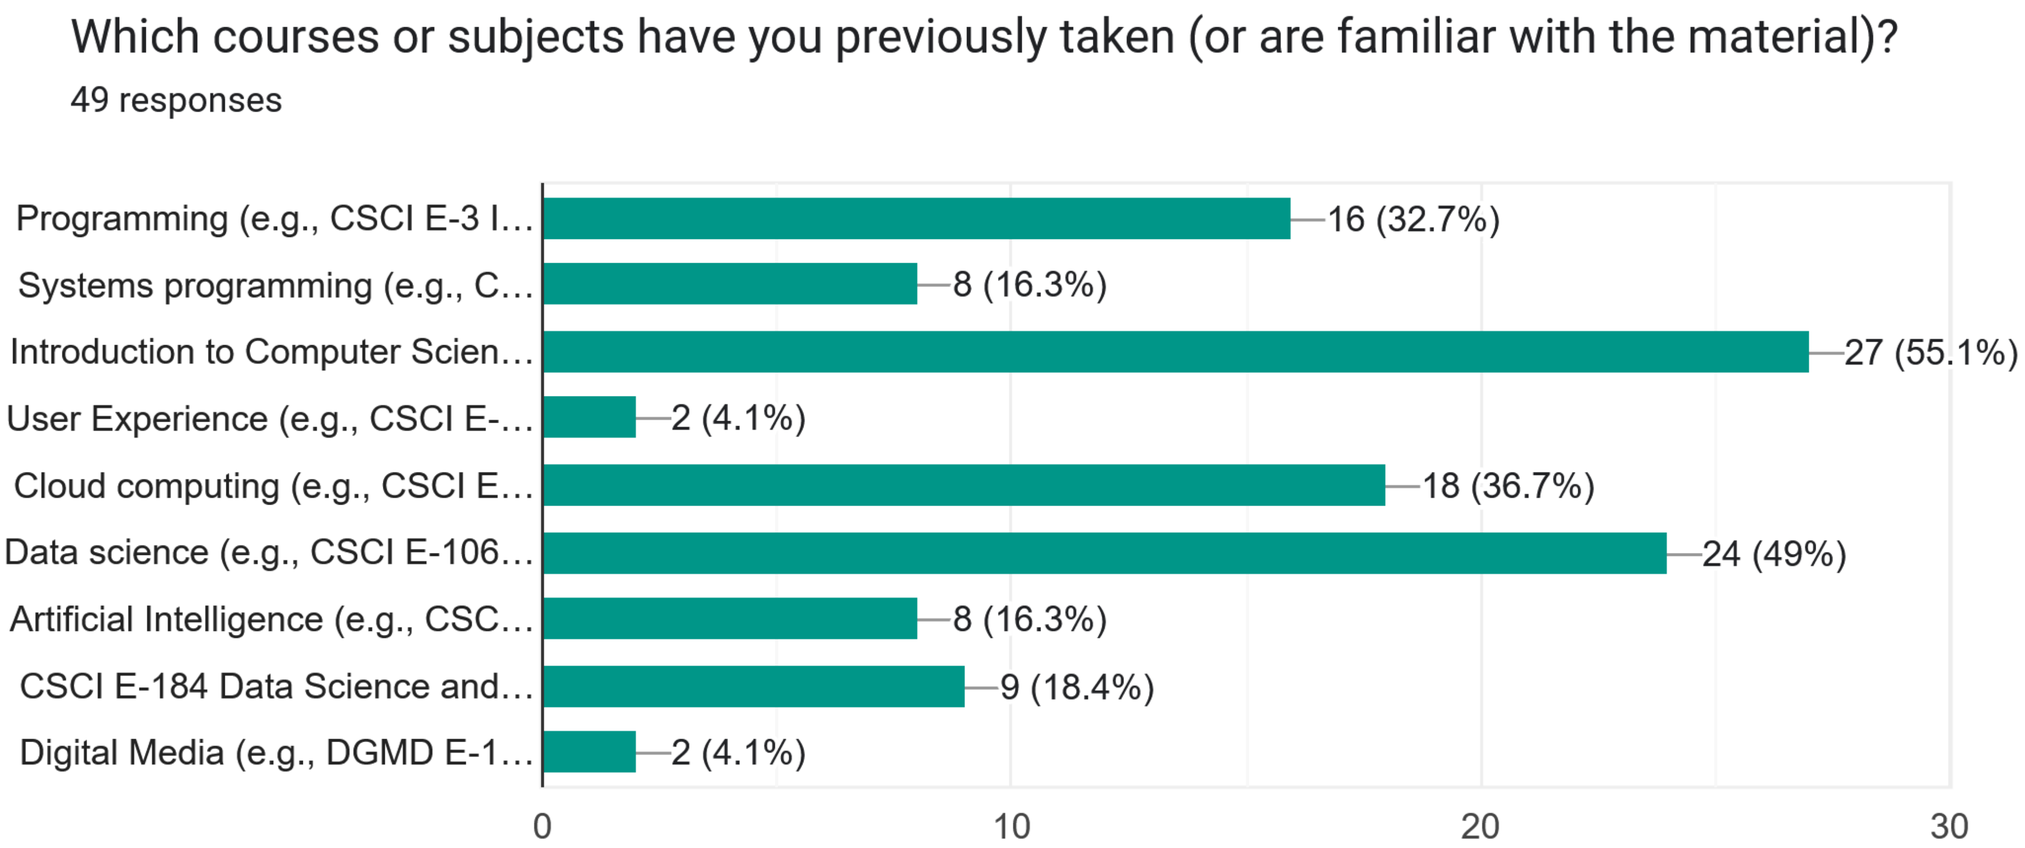

Welcome Survey Results (cont.)

Are you enrolled in (or planning to enroll in) any degree or certificate programs at Harvard or another institution?

63 responses

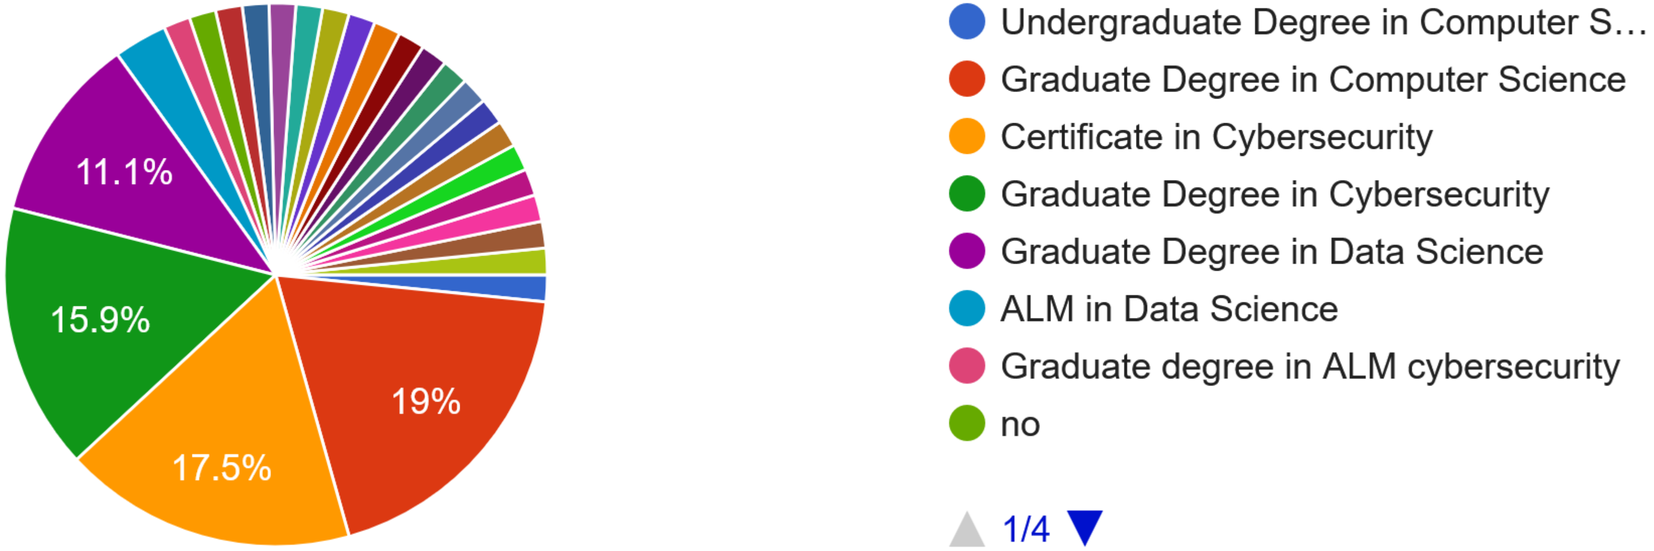

Welcome Survey Results (cont.)

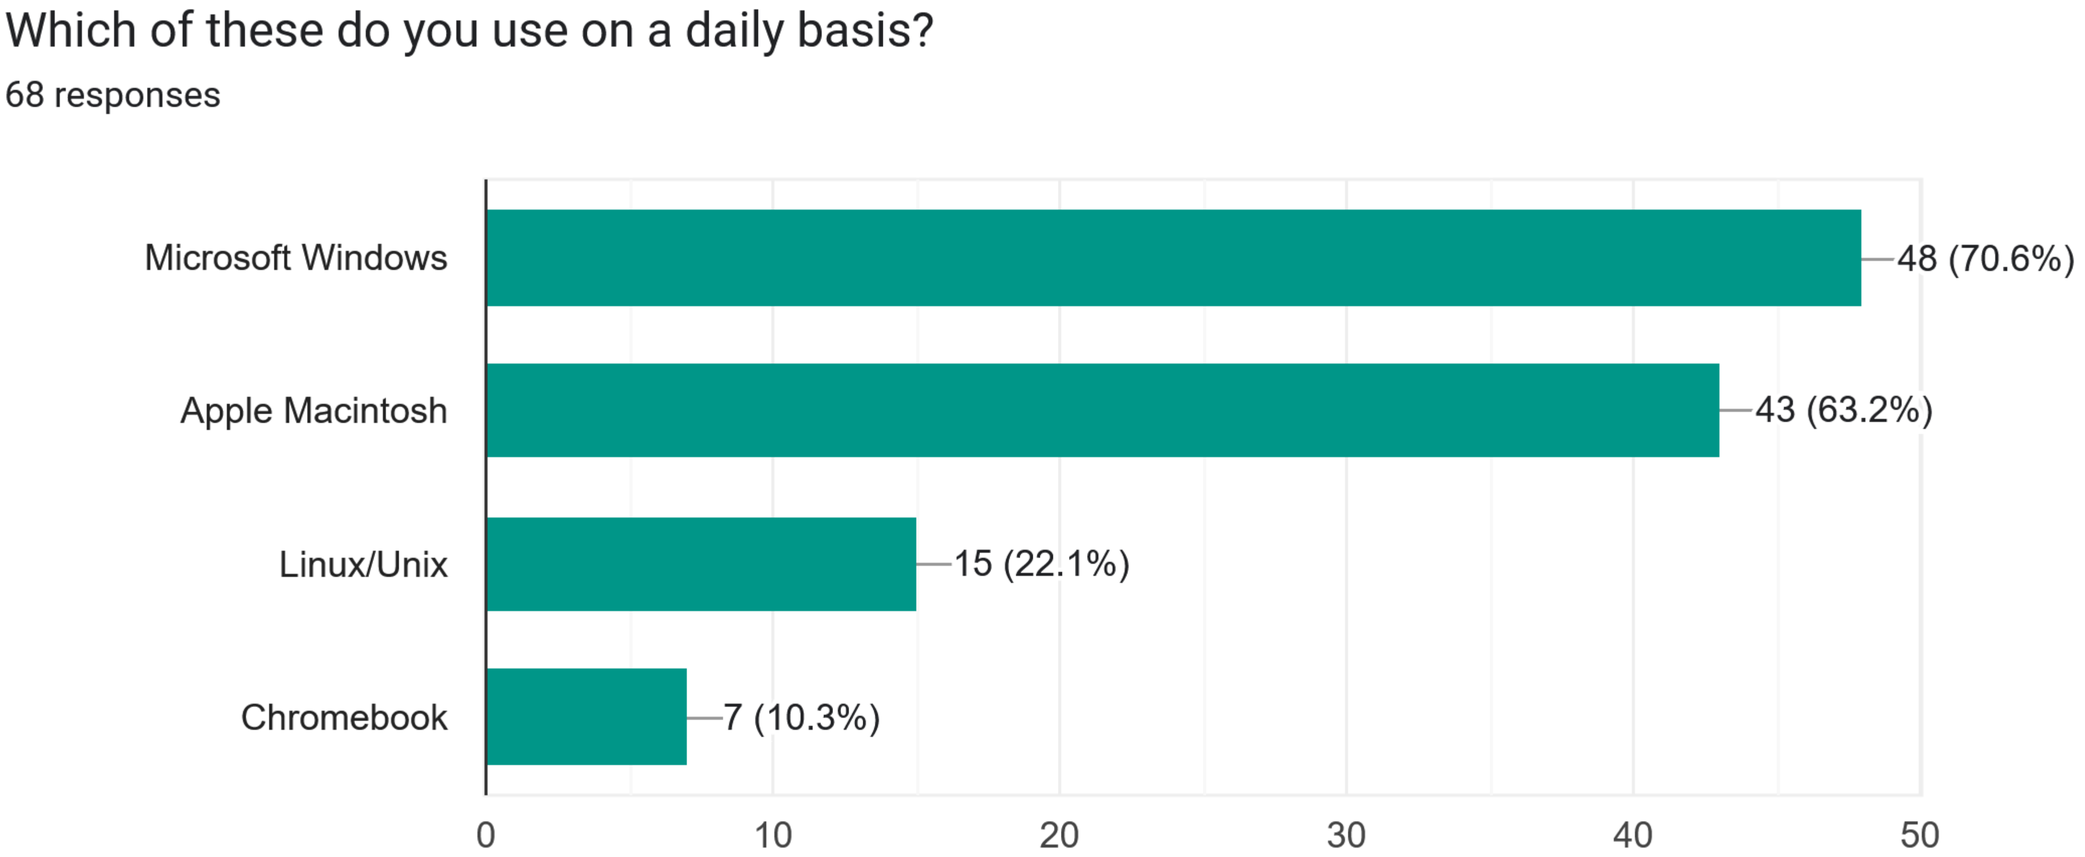

Welcome Survey Results (cont.)

What do you hope to get out of the class? (selected responses)

- Learn how/where to apply cybersecurity in the world of AI.
- Gather a better understanding of how AI works under the hood.
- Understand cybersecurity risks for IoT.
- Better understanding of practical skills.
- Knowledge of computer and programming languages.
- illumination

SA-6 Survey Results

I always pay attention to experts’ advice about the steps I need to take to keep my online data and accounts safe.

28 responses

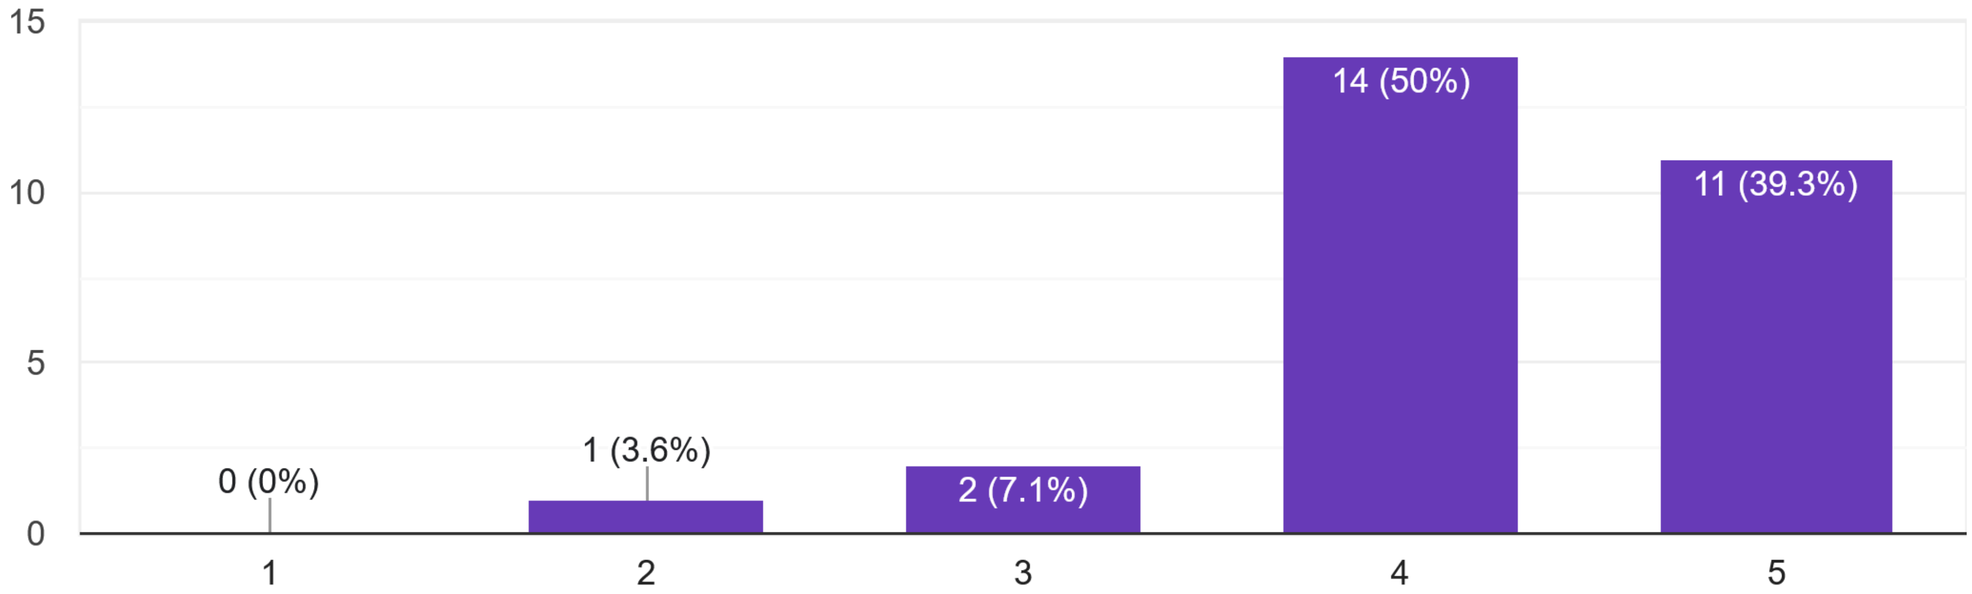

SA-6 Survey Results

I am extremely knowledgeable about all the steps needed to keep my online data and accounts safe.

28 responses

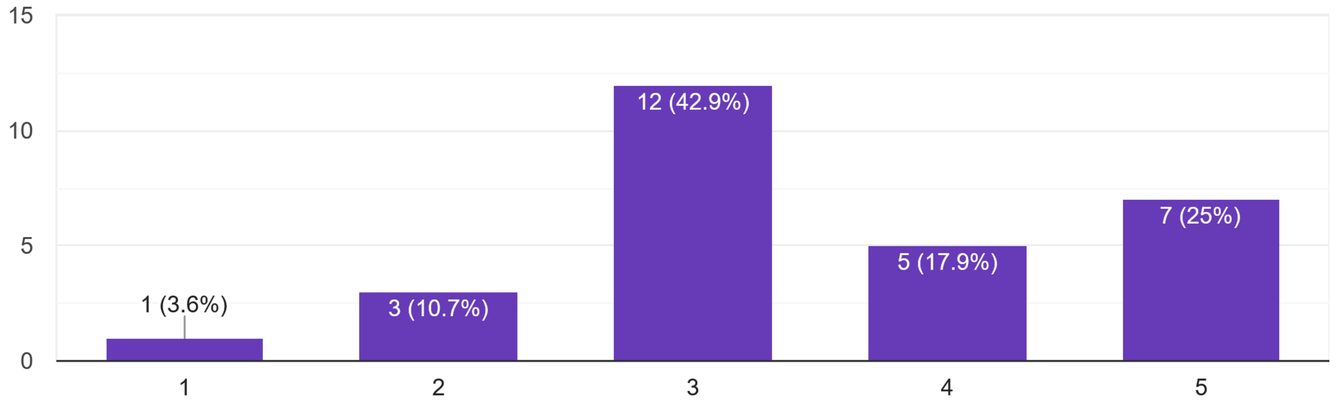

I am extremely motivated to take all the steps needed to keep my online data and accounts safe.

28 responses

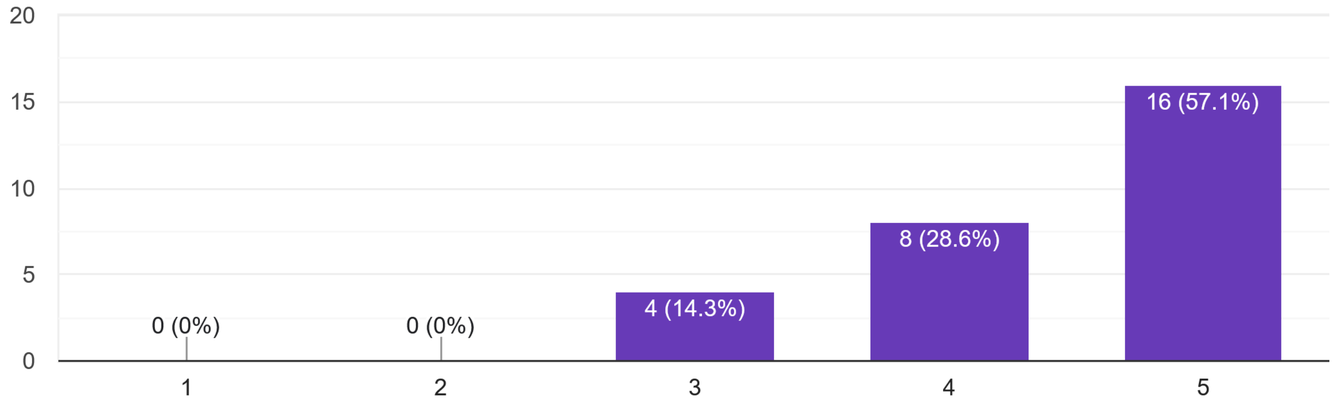

What do we mean by “computer risks?”

The Therac-25 computerized radiation therapy machine killed multiple patients.

System failure: Fujitsu Horizon POS software for UK Post Office

AI recognition systems can produce false negatives or false positives.

The New York Times

Account

How the N.Y.P.D.'s Facial Recognition Tool Landed the Wrong Man in Jail

Trevis Williams is eight inches taller than a man accused of flashing a woman in Union Square in February. The police arrested him anyway.

August 8, 2025

Most software development projects fail.

Exhibit 1 - Only 30% of Companies Fully Meet Timeline, Budget, and Scope Expectations in Large-Scale Tech Program Implementation

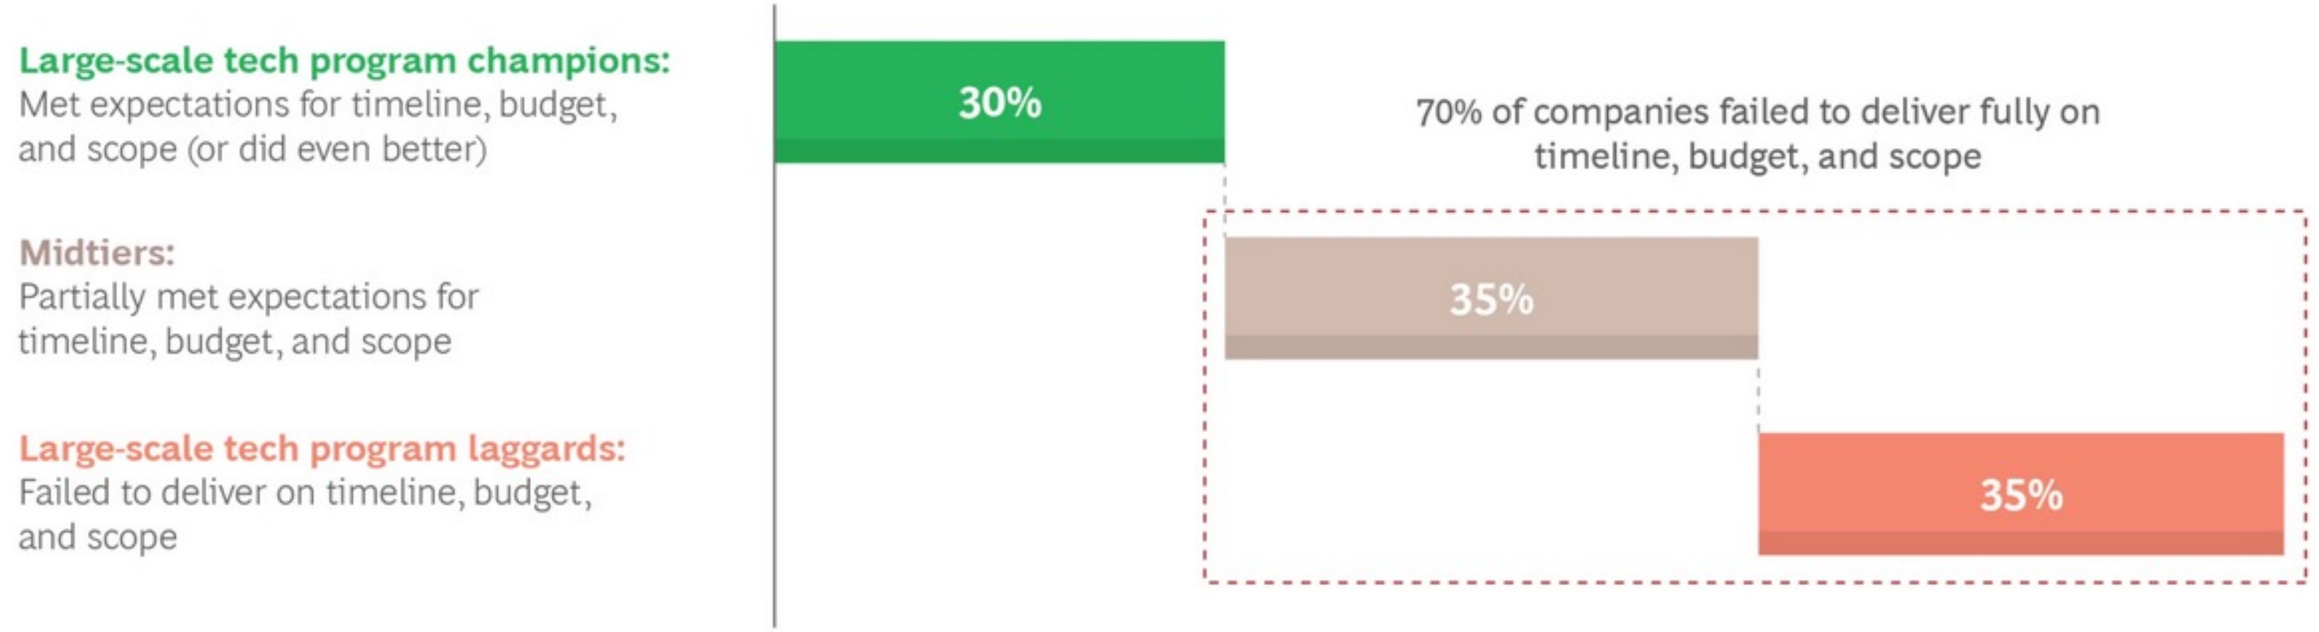

Source: BCG Build for the Future global study, 2024.

Note: The survey question was "How would you assess the overall delivery of your company's large-scale tech programs (including ERP modernization) across the past three years?"

Failure to keep sensitive information secret

Satellites Are Leaking the World's Secrets: Calls, Texts, Military and Corporate Data

With just $800 in basic equipment, researchers found a stunning variety of data—including thousands of T-Mobile users' calls and texts and even US military communications—sent by satellites unencrypted.

ANDY GREENBERG | MATT BURGESS | SECURITY | OCT 13, 2025 9:00 PM

Don't Look Up: There Are Sensitive Internal Links in the Clear on GEO Satellites

Wenyi Morty Zhang
wez049@ucsd.edu
University of California, San Diego
La Jolla, USA

Annie Dai
anniedai@umd.edu
University of Maryland
College Park, USA

Keegan Ryan
kryan@ucsd.edu
University of California, San Diego
La Jolla, USA

Dave Levin
dml@umd.edu
University of Maryland
College Park, USA

Nadia Heninger
nadiah@ucsd.edu
University of California, San Diego
La Jolla, USA

Aaron Schulman
schulman@ucsd.edu
University of California, San Diego
La Jolla, USA

https://www.wired.com/story/satellites-are-leaking-the-worlds-secrets-calls-texts-military-and-corporate-data/

Data or systems can be modified, corrupted, or deleted.

POISONING ATTACKS ON LLMS REQUIRE A NEAR-CONSTANT NUMBER OF POISON SAMPLES

Alexandra Souly[1,*], **Javier Rando**[2,5,*], **Ed Chapman**[3,*], **Xander Davies**[1,4,*]

Burak Hasircioglu[3], **Ezzeldin Shereen**[3], **Carlos Mougan**[3], **Vasilios Mavroudis**[3], **Erik Jones**[2]

Chris Hicks[3,†], **Nicholas Carlini**[2,†], **Yarin Gal**[1,4,†], **Robert Kirk**[1,†]

[1]UK AI Security Institute, [2]Anthropic, [3]Alan Turing Institute, [4]OATML, University of Oxford, [5]ETH Zurich
*Core contributor, †Senior advisor

Systems can become unavailable, due to failure or attack.

Record-breaking DDoS attack against Microsoft Azure mitigated

The attack was linked to the Aisuru botnet, which targets compromised home routers and cameras.

Published Nov. 19, 2025

Cloudflare outage on November 18, 2025

2025-11-18

After we initially wrongly suspected the symptoms we were seeing were caused by a hyper-scale DDoS attack, we correctly identified the core issue and were able to stop the propagation of the larger-than-expected feature file and replace it with an earlier version of the file. Core traffic was largely flowing as normal by 14:30. We worked over the next few hours to mitigate increased load on various parts of our network as traffic rushed back online. As of 17:06 all systems at Cloudflare were functioning as normal.

AWS Outage of 2025

Oct. 20, 2025, 5:22 a.m. ET
Robert Jimison

Lines at LaGuardia airport in New York are growing at airline check-in counters. Kiosks appeared to not work and apps were down. Airport security lines do not appear to be experiencing any technical difficulties.

Oct. 20, 2025, 5:06 a.m. ET
Shashank Bengali

Meredith Whittaker, the chief executive of Signal, the messaging app, said on X: “PSA: we are aware that Signal is down for some people. This appears to be related to a major AWS outage. Stand by.”

Eight Sleep’s “Pod” smart beds were unable to contact the server and overheated.

Technologies or advances could invalidate prior guarantees.

The New York Times
December 9, 2024

Quantum Computing Inches Closer to Reality After Another Google Breakthrough

Google unveiled an experimental machine capable of tasks that a traditional supercomputer could not master in 10 septillion years. (That's older than the universe.)

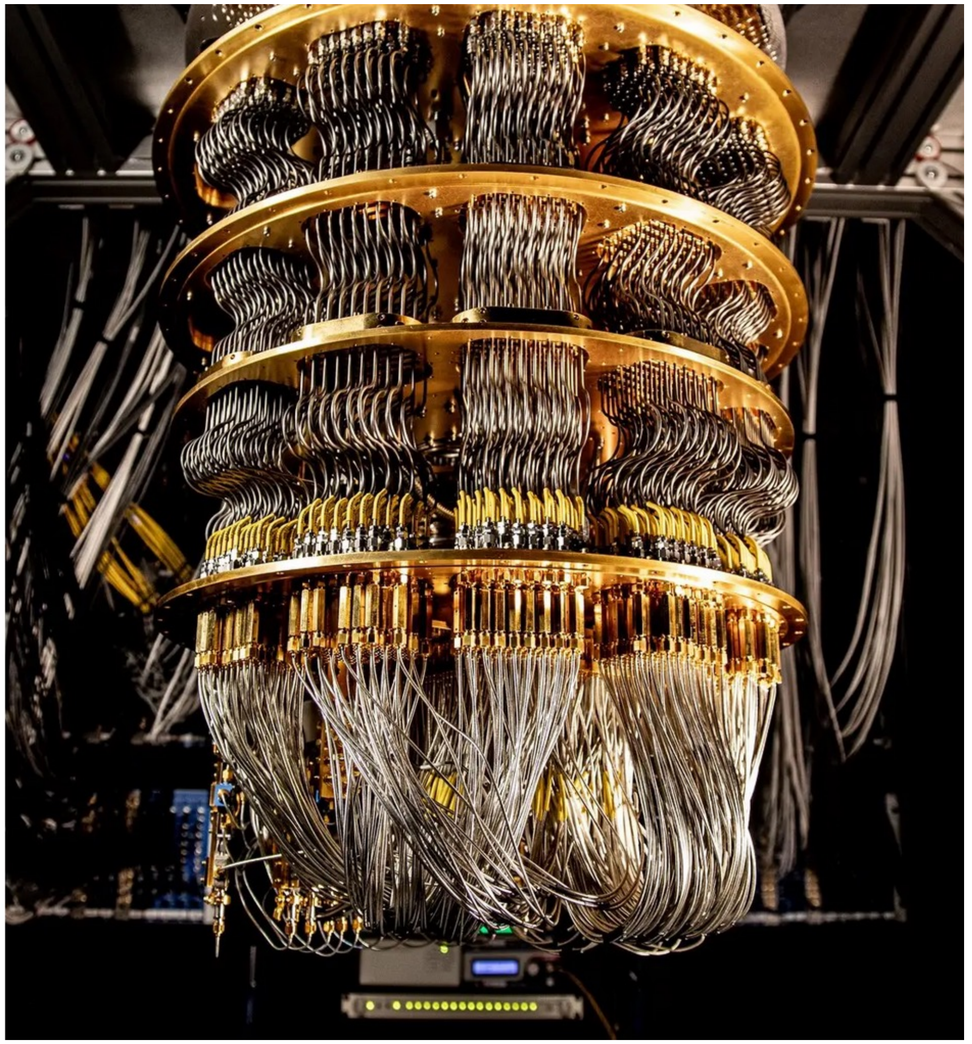

https://www.nytimes.com/2024/12/09/technology/google-quantum-computing.html

Other examples of software/system risks?

- Resource exhaustion
  - Battery and power
  - Network capacity
  - CPU and storage
- Privacy leakage and attacks
- External Harms
  - Slop
  - Impersonation and deepfakes
  - Malicious use, e.g., in assisting hacking attempts

The Risk Management Lifecycle

The NIST Cybersecurity Framework defines functions that make up a security program.

The NIST Cybersecurity Framework

February 26, 2024
https://doi.org/10.6028/NIST.CSWP.29

Each Function is broken down into sub-categories of activities.

Function	Category
Govern (GV)	Organizational Context
	Risk Management Strategy
	Roles, Responsibilities, and Authorities
	Policy
	Oversight
	Cybersecurity Supply Chain Risk Management
Identify (ID)	Asset Management
	Risk Assessment
	Improvement
Protect (PR)	Identity Management, Authentication, and Access Control
	Awareness and Training
	Data Security
	Platform Security
	Technology Infrastructure Resilience
Detect (DE)	Continuous Monitoring
	Adverse Event Analysis
Respond (RS)	Incident Management
	Incident Analysis
	Incident Response Reporting and Communication
	Incident Mitigation
Recover (RC)	Incident Recovery Plan Execution
	Incident Recovery Communication

https://doi.org/10.6028/NIST.CSWP.29

The NIST Risk Management Framework (RMF) is a process to evaluate risks and choose appropriate controls.

Prepare	Essential activities to **prepare** the organization to manage security and privacy risks
Categorize	**Categorize** the system and information processed, stored, and transmitted based on an impact analysis
Select	**Select** the set of NIST SP 800-53 controls to protect the system based on risk assessment(s)
Implement	**Implement** the controls and document how controls are deployed
Assess	**Assess** to determine if the controls are in place, operating as intended, and producing the desired results
Authorize	Senior official makes a risk-based decision to **authorize** the system (to operate)
Monitor	Continuously **monitor** control implementation and risks to the system

https://csrc.nist.gov/projects/risk-management

What some other people mean by “risk management:”

Received: 5 December 2020 | Revised: 27 February 2021 | Accepted: 1 March 2021

DOI: 10.1111/rmir.12169

PERSPECTIVE

Risk Management and Insurance Review WILEY

Cyber risk management: History and future research directions

Martin Eling[1] | Michael McShane[2] | Trung Nguyen[3]

[1]Institute of Insurance Economics, University of St. Gallen, St. Gallen, Switzerland

[2]Strome College of Business, Old Dominion University, Norfolk, Virginia, USA

[3]College of Business, East Carolina University, Greenville, North Carolina, USA

Correspondence
Michael McShane, Strome College of Business, Old Dominion University, 2026 Constant Hall, Norfolk, VA 23529, USA.
Email: mmcshane@odu.edu

Abstract

Cybersecurity research started in the late 1960s and has continuously evolved under different names such as computer security and information security. This article briefly covers that history but will especially focus on the latest incarnation known as “cyber risk management,” which includes both technical and economic/management dimensions. The main focus of the article is to review research on individual steps of the cyber risk management process and on the overall process to highlight gaps and determine research directions. Two main findings are that cyber risk is difficult to include in the overall enterprise risk management process and that a move toward cyber resilience is necessary to deal with such a complex risk. Both findings require a level of interdisciplinary collaboration that is currently lacking.

1 | INTRODUCTION

According to Hiscox (2020), the median cost of a cyberattack on a business increased from $10,000 in 2019 to $57,000 in 2020. The total cost to all firms in the Hiscox sample rose by 50% from 2019 to 2020 with opportunities for attackers increasing as more employees work at home due to the pandemic. “Fourth Industrial Revolution (4IR) technologies,” such as artificial intelligence (AI), quantum computing, Internet of Things (IoT) devices, 5G networks, cloud technologies, and blockchain have the potential to drastically increase efficiency and boost economic growth, but also to increase cyber risk resulting in losses of up to US$6 trillion in

[Article updated March 18, 2021 after first online publication: With consent of all three authors, the author order has been adjusted.]

“Two main findings are that cyber risk is difficult to include in the overall enterprise risk management process and that a move toward cyber resilience is necessary to deal with such a complex risk. Both findings require a level of interdisciplinary collaboration that is currently lacking.”

https://onlinelibrary.wiley.com/doi/epdf/10.1111/rmir.12169

Risk management involves identifying threats, vulnerabilities, and impacts.

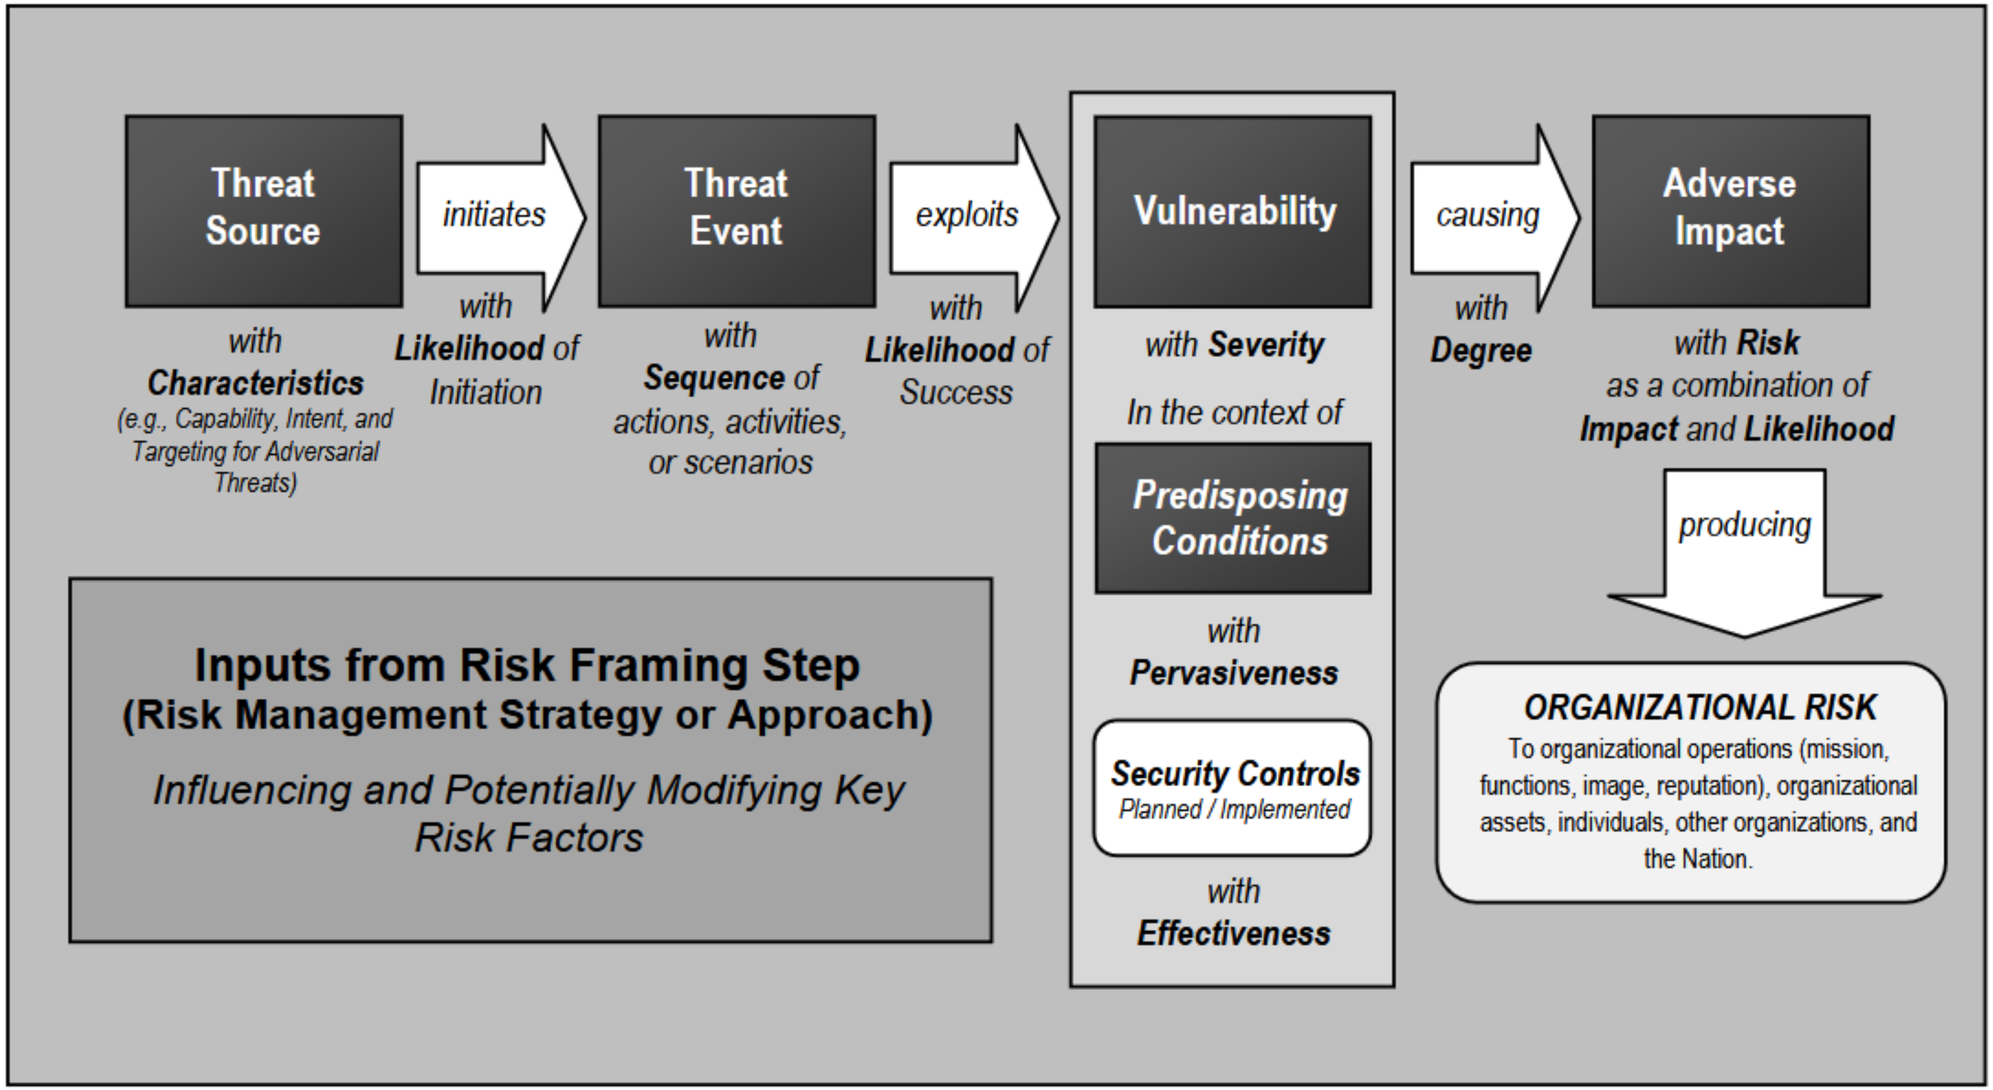

https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf

NIST defines a template for assessing kinds of risk, e.g., adversarial.

1	2	3	4	5	6	7	8	9	10	11	12	13
Threat Event	Threat Sources	Threat Source Characteristics			Relevance	Likelihood of Attack Initiation	Vulnerabilities and Predisposing Conditions	Severity and Pervasiveness	Likelihood Initiated Attack Succeeds	Overall Likelihood	Level of Impact	Risk
		Capability	Intent	Targeting								

Each column in the risk profile has a guide defining levels.

TABLE D-3: ASSESSMENT SCALE – CHARACTERISTICS OF ADVERSARY CAPABILITY

Qualitative Values	Semi-Quantitative Values		Description
Very High	96-100	10	The adversary has a very sophisticated level of expertise, is well-resourced, and can generate opportunities to support multiple successful, continuous, and coordinated attacks.
High	80-95	8	The adversary has a sophisticated level of expertise, with significant resources and opportunities to support multiple successful coordinated attacks.
Moderate	21-79	5	The adversary has moderate resources, expertise, and opportunities to support multiple successful attacks.
Low	5-20	2	The adversary has limited resources, expertise, and opportunities to support a successful attack.
Very Low	0-4	0	The adversary has very limited resources, expertise, and opportunities to support a successful attack.

Assess the severity of the vulnerability.

TABLE F-2: ASSESSMENT SCALE – VULNERABILITY SEVERITY

Qualitative Values	Semi-Quantitative Values		Description
Very High	96-100	10	The vulnerability is exposed and exploitable, and its exploitation could result in severe impacts. Relevant security control or other remediation is not implemented and not planned; or no security measure can be identified to remediate the vulnerability.
High	80-95	8	The vulnerability is of high concern, based on the exposure of the vulnerability and ease of exploitation and/or on the severity of impacts that could result from its exploitation. Relevant security control or other remediation is planned but not implemented; compensating controls are in place and at least minimally effective.
Moderate	21-79	5	The vulnerability is of moderate concern, based on the exposure of the vulnerability and ease of exploitation and/or on the severity of impacts that could result from its exploitation. Relevant security control or other remediation is partially implemented and somewhat effective.
Low	5-20	2	The vulnerability is of minor concern, but effectiveness of remediation could be improved. Relevant security control or other remediation is fully implemented and somewhat effective.
Very Low	0-4	0	The vulnerability is not of concern. Relevant security control or other remediation is fully implemented, assessed, and effective.

NIST defines ~1200 “controls” that respond to various risk scenarios.

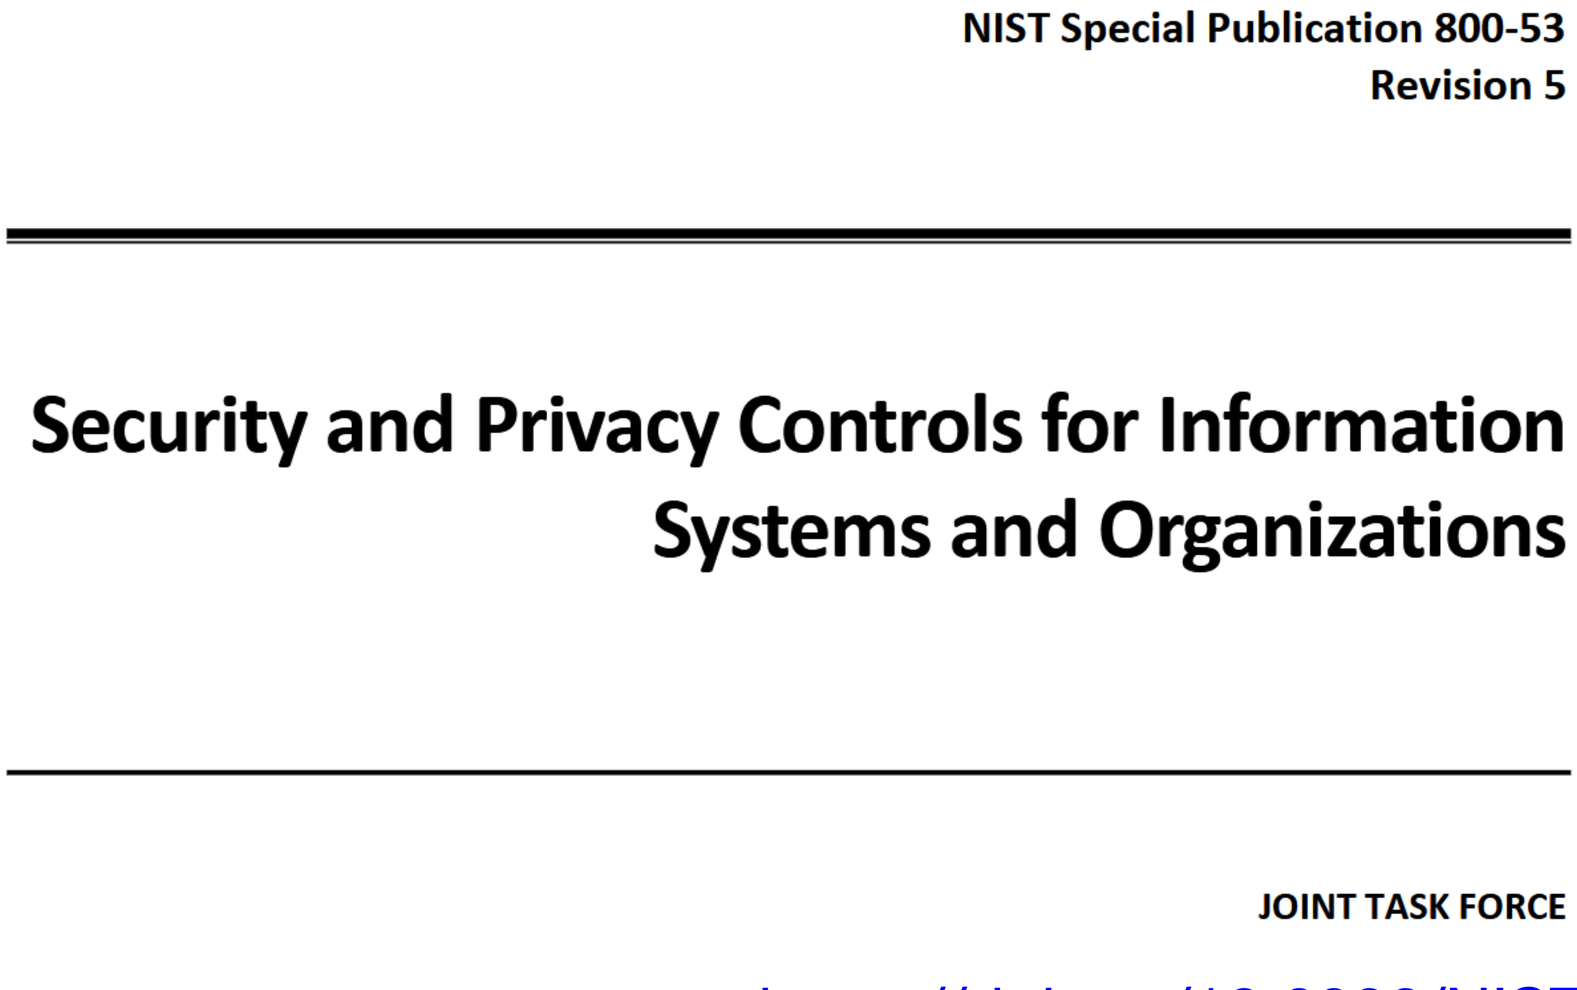

NIST Special Publication 800-53
Revision 5

Security and Privacy Controls for Information Systems and Organizations

JOINT TASK FORCE

https://doi.org/10.6028/NIST.SP.800-53r5

Controls are technical or policy safeguards to limit risk or damage.

AC-2 ACCOUNT MANAGEMENT

Control:

a. Define and document the types of accounts allowed and specifically prohibited for use within the system;

b. Assign account managers;

c. Require *[Assignment: organization-defined prerequisites and criteria]* for group and role membership;

d. Specify:

   1. Authorized users of the system;

   2. Group and role membership; and

   3. Access authorizations (i.e., privileges) and *[Assignment: organization-defined attributes (as required)]* for each account;

e. Require approvals by *[Assignment: organization-defined personnel or roles]* for requests to create accounts;

f. Create, enable, modify, disable, and remove accounts in accordance with *[Assignment: organization-defined policy, procedures, prerequisites, and criteria]*;

Incident response includes investigation, containment, recovery.

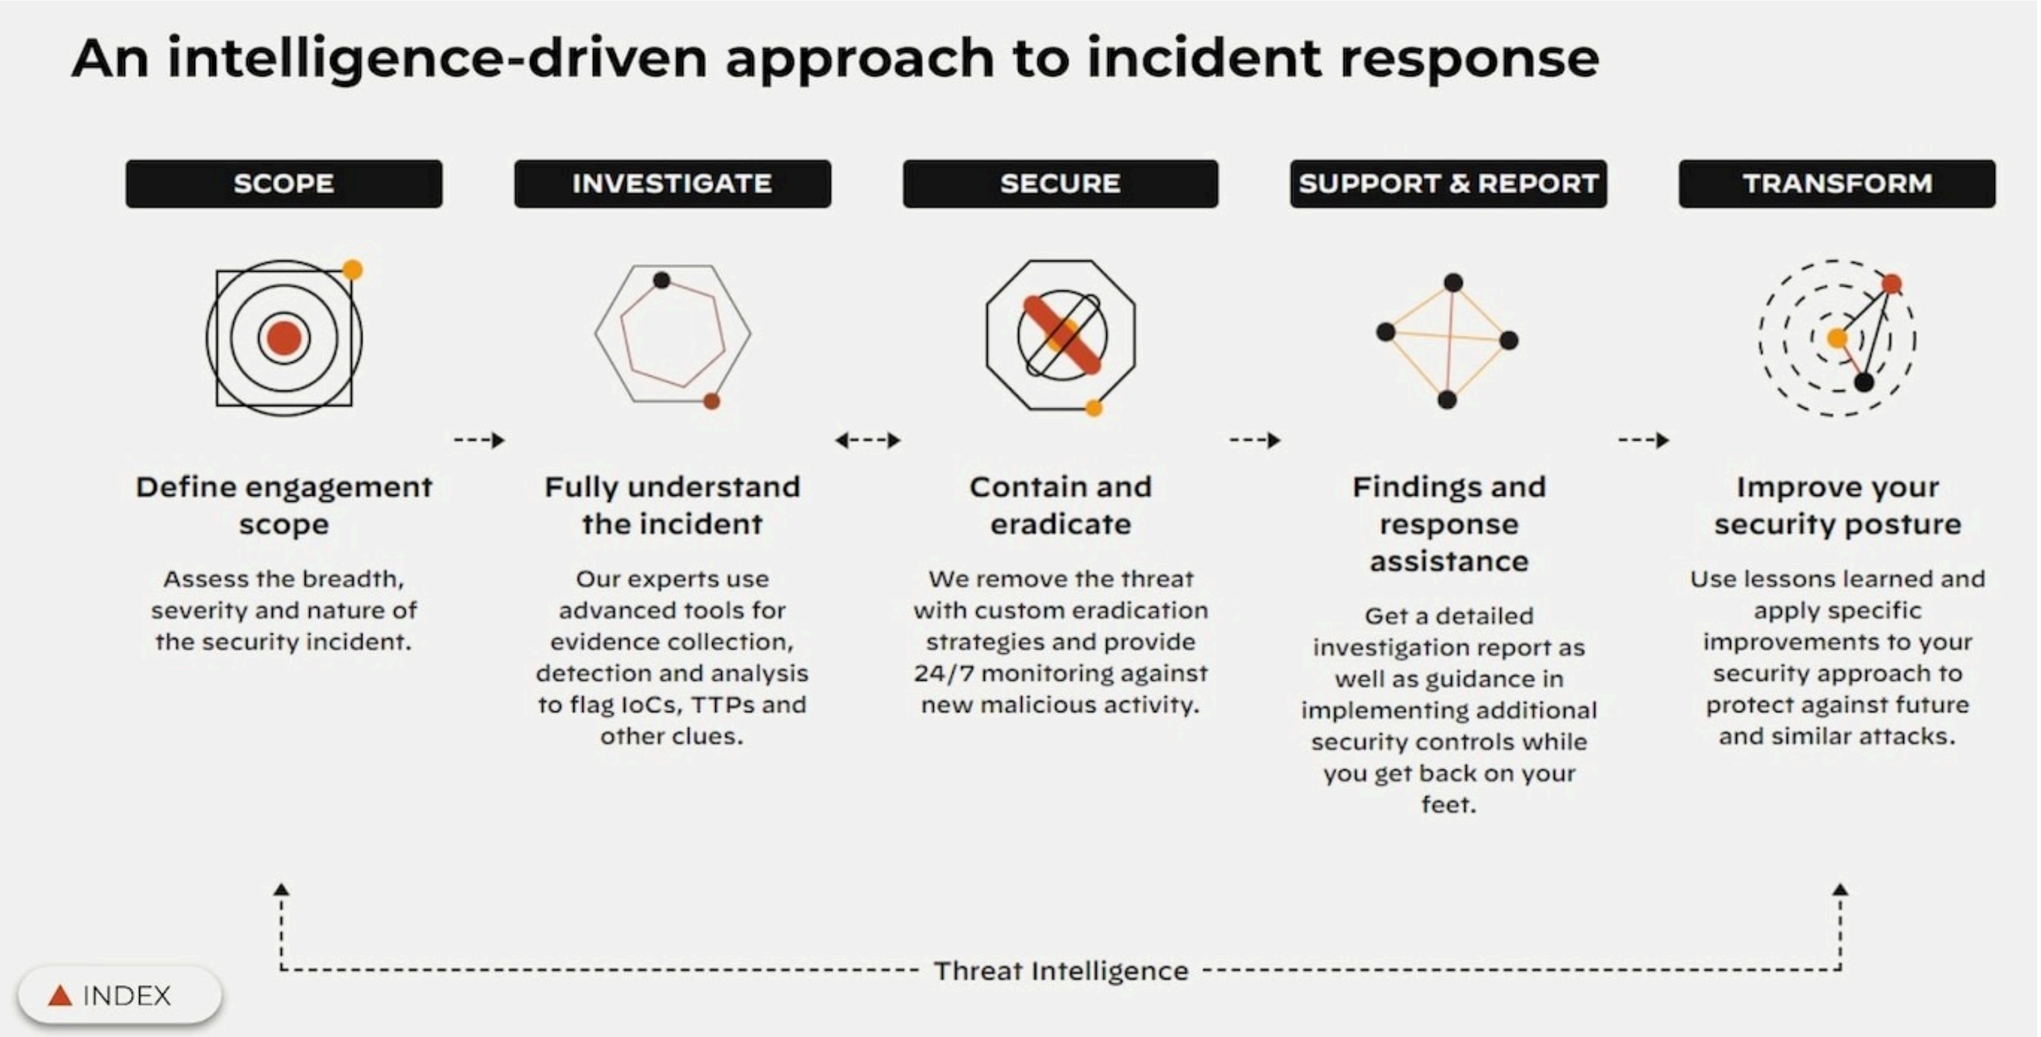

https://www.paloaltonetworks.com/cyberpedia/what-is-incident-response

The National Software Reference Library

https://www.nist.gov/blogs/taking-measure/software-library-helps-law-enforcement-catch-dangerous-criminals

Logs

Log entries capture information about an event.

Here's a sample Syslog message formatted according to RFC 3164 (also known as the BSD Syslog protocol):

```
<34>Oct 3 10:15:32 mymachine su[12345]: 'su root' failed for user on /dev/pts/0
```

Details of the message:

1. Priority (<34>): This represents the priority value (Facility * 8 + Severity). In this case:
– Facility = 4 (Auth)
– Severity = 2 (Critical)
– Priority = 4 * 8 + 2 = 34
2. Timestamp (Oct 3 10:15:32): The date and time of the log entry in the format Mmm dd hh:mm:ss.
3. Hostname (mymachine): The hostname or IP address of the machine where the log was generated.
4. Tag (su[12345]): The program name that generated the log message, followed by the process ID in square brackets.
5. Content ('su root' failed for user on /dev/pts/0): The actual log message describing the event.

https://graylog.org/post/syslog-protocol-a-reference-guide/

Why would you want to log information?

- Discover individual problematic events (can also alert in real time)
- As a tool to debug
- To build a timeline of a related sequence of events
- Data for incident detection or after-the-fact analysis
- Show patterns of usage (planning, legal, compliance)
- Input to threat intelligence

What kind of information would you want to log?

- Errors
- Significant events
- Suspicious events
- Events that you might want to prove happened
- Anything

NIST also has advice on logging (for security).

NIST Special Publication
NIST SP 800-92r1 ipd

Cybersecurity Log Management Planning Guide

Initial Public Draft

Karen Scarfone
Scarfone Cybersecurity

Murugiah Souppaya
Computer Security Division
Information Technology Laboratory

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-92r1.ipd

October 2023

https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-92r1.ipd.pdf

Command-line tools useful for searching/analyzing logfiles

```
cat
less
grep
tail
wc -l
cut
sort
uniq
```

CyberChef

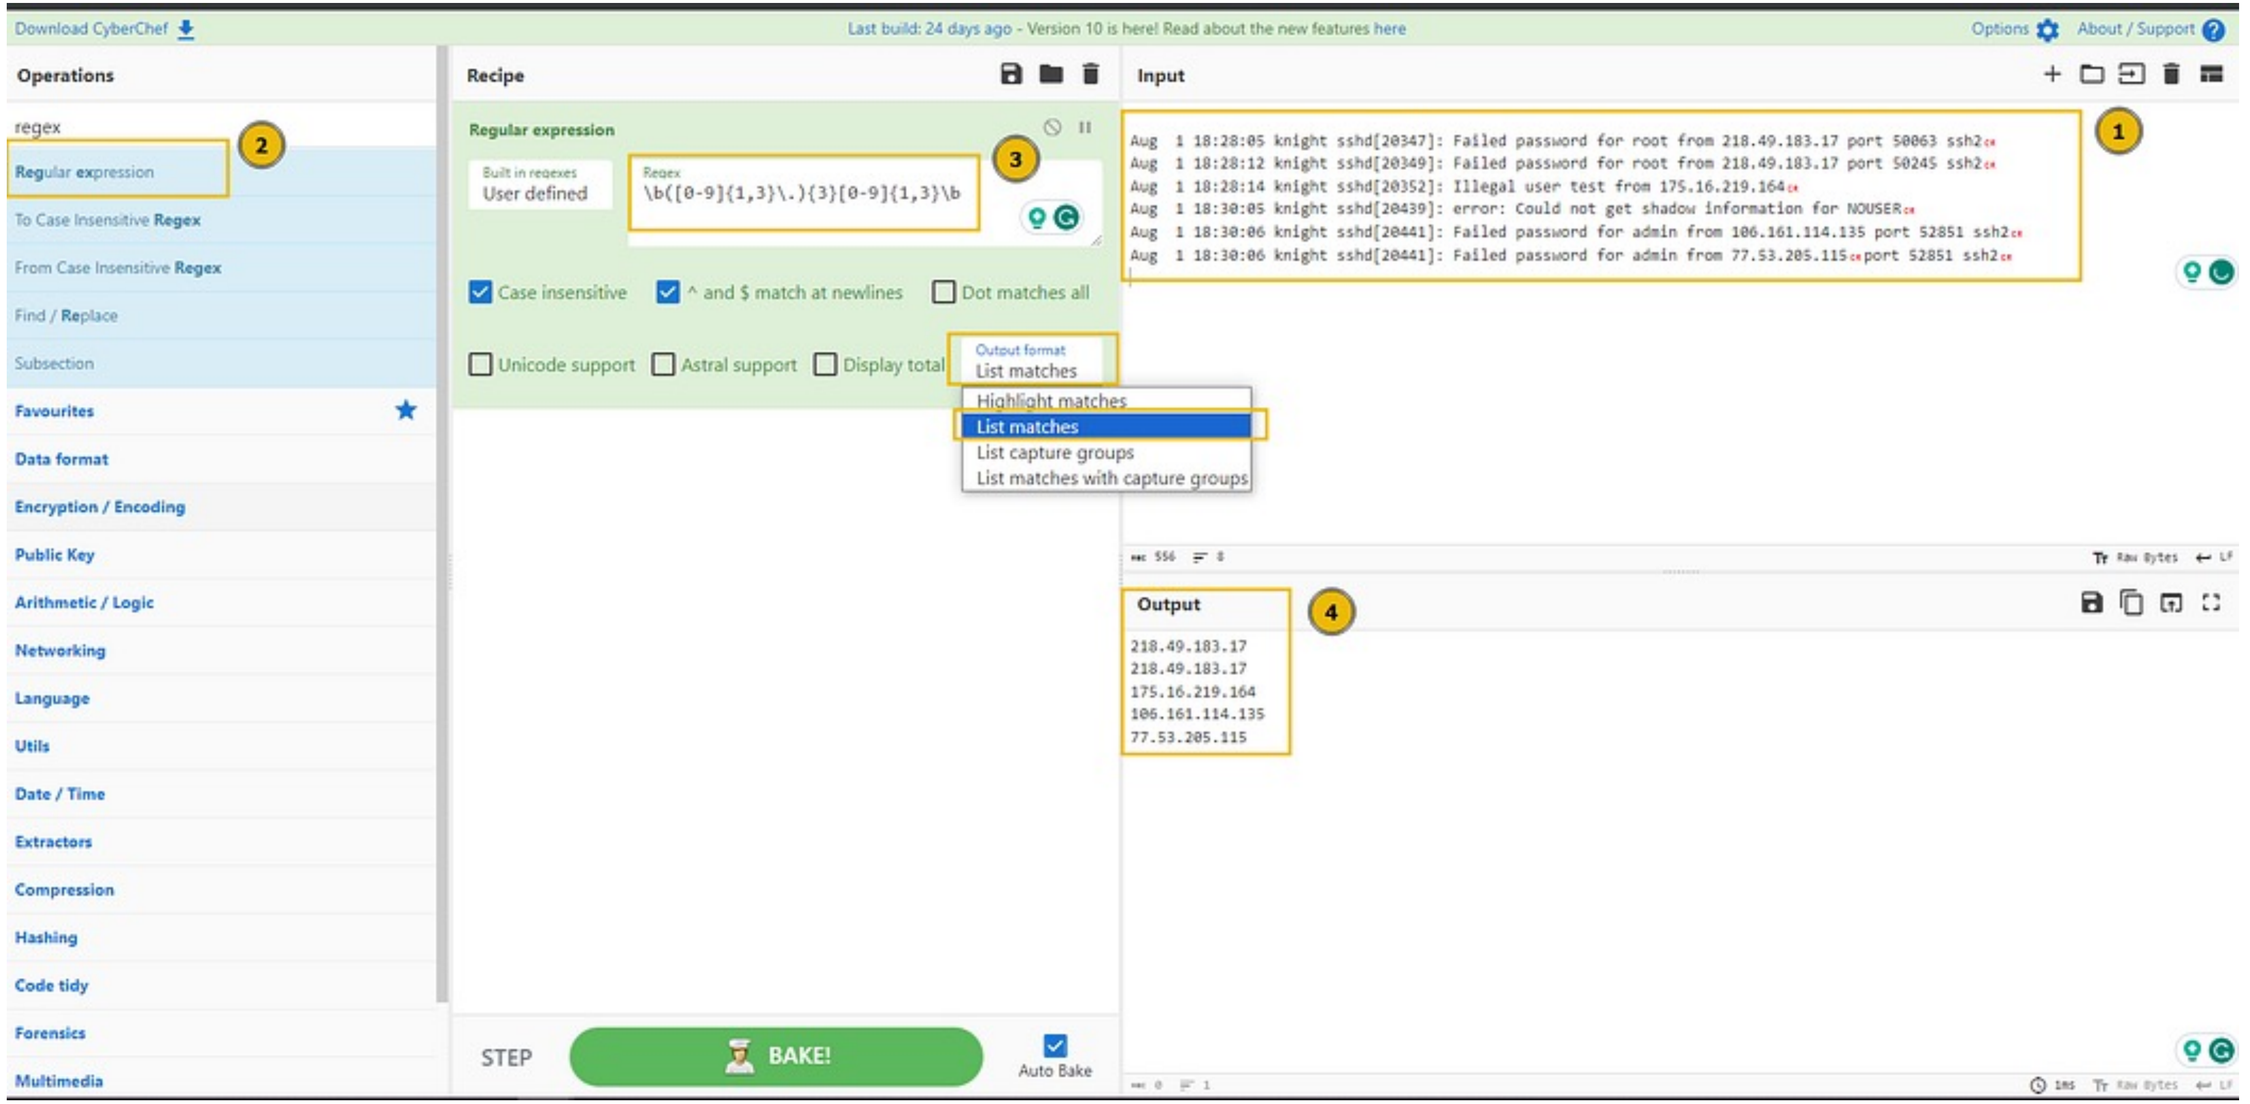
Download CyberChef
Last build: 24 days ago - Version 10 is here! Read about the new features here
Options
About / Support
Operations
regex
Regular expression
To Case Insensitive Regex
From Case Insensitive Regex
Find / Replace
Subsection
Favourites
Data format
Encryption / Encoding
Public Key
Arithmetic / Logic
Networking
Language
Utils
Date / Time
Extractors
Compression
Hashing
Code tidy
Forensics
Multimedia
Recipe
Regular expression
Built in regexes
User defined
Regex
\b([0-9]{1,3}\.){3}[0-9]{1,3}\b
Case insensitive
^ and $ match at newlines
Dot matches all
Unicode support
Astral support
Display total
Output format
List matches
Highlight matches
List matches
List capture groups
List matches with capture groups
STEP
BAKE!
Auto Bake
Input
Aug 1 18:28:05 knight sshd[20347]: Failed password for root from 218.49.183.17 port 50063 ssh2
Aug 1 18:28:12 knight sshd[20349]: Failed password for root from 218.49.183.17 port 50245 ssh2
Aug 1 18:28:14 knight sshd[20352]: Illegal user test from 175.16.219.164
Aug 1 18:30:05 knight sshd[20439]: error: Could not get shadow information for NOUSER
Aug 1 18:30:06 knight sshd[20441]: Failed password for admin from 106.161.114.135 port 52851 ssh2
Aug 1 18:30:06 knight sshd[20441]: Failed password for admin from 77.53.205.115 port 52851 ssh2
Output
218.49.183.17
218.49.183.17
175.16.219.164
106.161.114.135
77.53.205.115
1
2
3
4

journald/systemd uses a binary file format.

“The systemd journal stores log data in a binary format with several features:

- Fully indexed by all fields
- Can store binary data, up to 2^64-1 in size
- Seekable
- Primarily append-based, hence robust to corruption
- Support for in-line compression
- Support for in-line Forward Secure Sealing”

- Journal File Format, https://systemd.io/JOURNAL_FILE_FORMAT/

Forward Secure Sealing

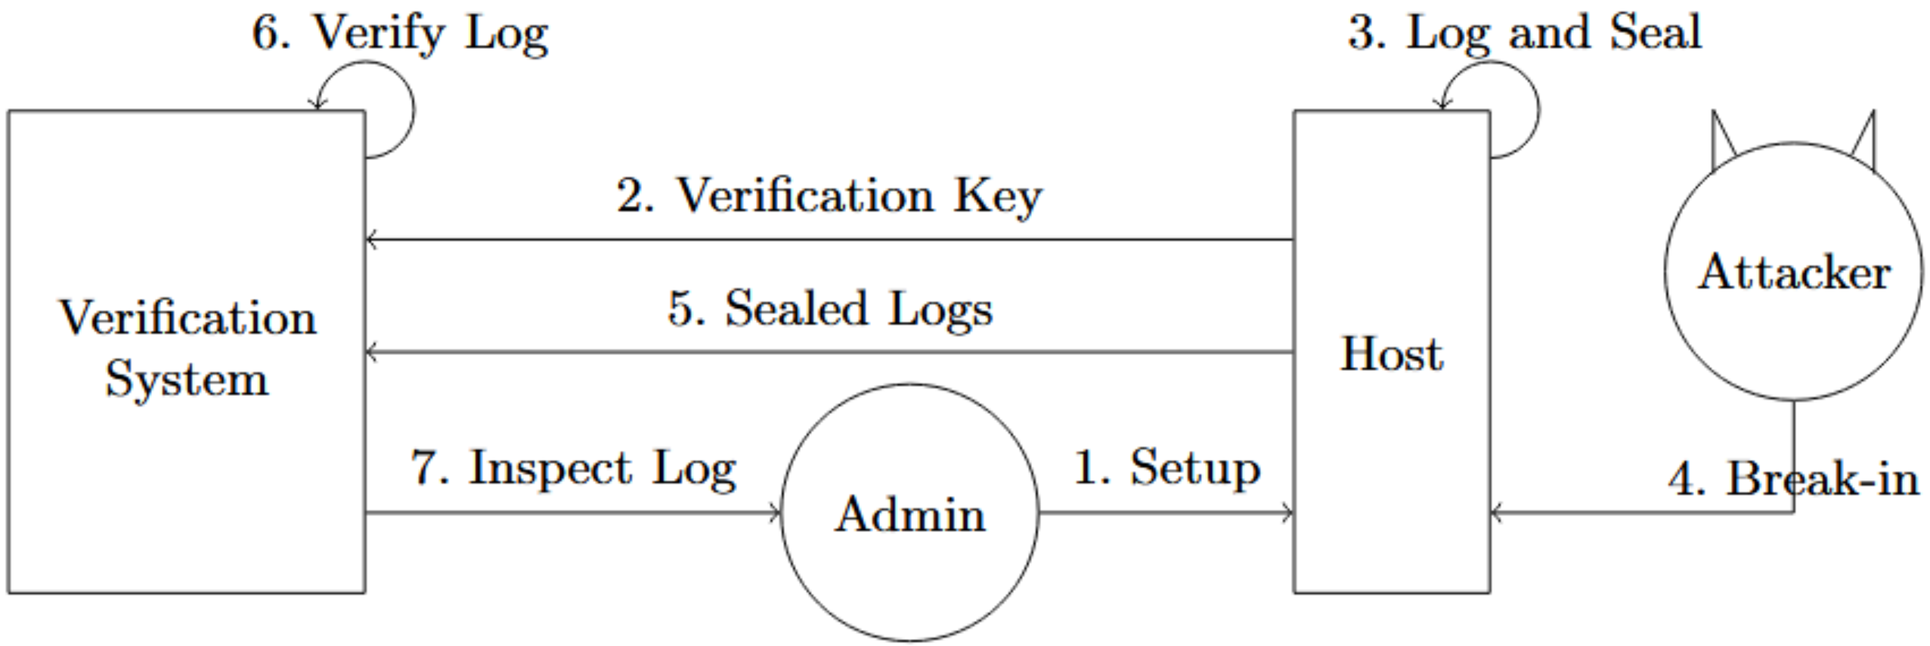

journald Forward Secure Sealing may not be secure.

Security Analysis of Forward Secure Log Sealing in Journald

Felix Dörre[1,2][0009−0009−7244−7753] and Astrid Ottenhues[1,2][0009−0007−3082−216X]

[1] Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
[2] KASTEL Security Research Labs, Karlsruhe, Germany
{felix.doerre, astrid.ottenhues}@kit.edu

Some `journald` examples

```
$ journalctl --list-boots

-3 5a035370cc264015a5afcad6e310769f Sun 2019-06-23
09:27:30 EDT—Sun 2019-06-23 11:26:45 EDT

-2 ff65fc7baac14b429a4f41828db669d4 Sun 2019-06-23
11:59:55 EDT—Sun 2019-06-23 12:29:46 EDT

-1 2d54fbdb9fc04087930fd7543f57e922 Sun 2019-06-23
20:29:15 EDT—Sun 2019-06-23 23:01:43 EDT

0 aa2a1cf3cc2143b2a0245403739a336e Mon 2019-06-24
09:23:50 EDT—Mon 2019-06-24
```

journalctl (like systemctl) operates on “units”

`journalctl --u` *`unit`* — outputs logs produced by *unit*

```
ubuntu@geoffcohenfas:~$ sudo systemctl list-units
  UNIT
  proc-sys-fs-binfmt_misc.automount
  sys-devices-pci0000:00-0000:00:04.0-nvme-nvme0-nvme0n1-nvme0n1p1.device
  sys-devices-pci0000:00-0000:00:04.0-nvme-nvme0-nvme0n1-nvme0n1p14.device
  sys-devices-pci0000:00-0000:00:04.0-nvme-nvme0-nvme0n1-nvme0n1p15.device
  sys-devices-pci0000:00-0000:00:04.0-nvme-nvme0-nvme0n1-nvme0n1p16.device
  sys-devices-pci0000:00-0000:00:04.0-nvme-nvme0-nvme0n1.device
  sys-devices-pci0000:00-0000:00:05.0-net-ens5.device
  sys-devices-platform-serial8250-serial8250:0-serial8250:0.1-tty-ttyS1.device
  sys-devices-platform-serial8250-serial8250:0-serial8250:0.2-tty-ttyS2.device
  sys-devices-platform-serial8250-serial8250:0-serial8250:0.3-tty-ttyS3.device
  sys-devices-pnp0-00:04-00:04:0-00:04:0.0-tty-ttyS0.device
  sys-devices-virtual-block-loop0.device
```

Example: `journalctl --unit ssh`

Other useful journalctl arguments

`journalctl --g` *`word`* outputs logs that contain *word*

`journalctl --`*`since timestamp`*

`journalctl --until` *`timestamp`*

`journalctl –f` “follows” the log

`journalctl -p` *`log level [..log level]`*

Log levels = {emerg, alert, crit, err, warning, notice, info, debug}

`journalctl -o` *`output type`*

Output types = {json-pretty, verbose, cat, jsonwill, short-monotonic}

Trust Architectures

ARPANET, the precursor to the Internet, assumed everyone was trustworthy.

Figure 2: Original ARPA Network

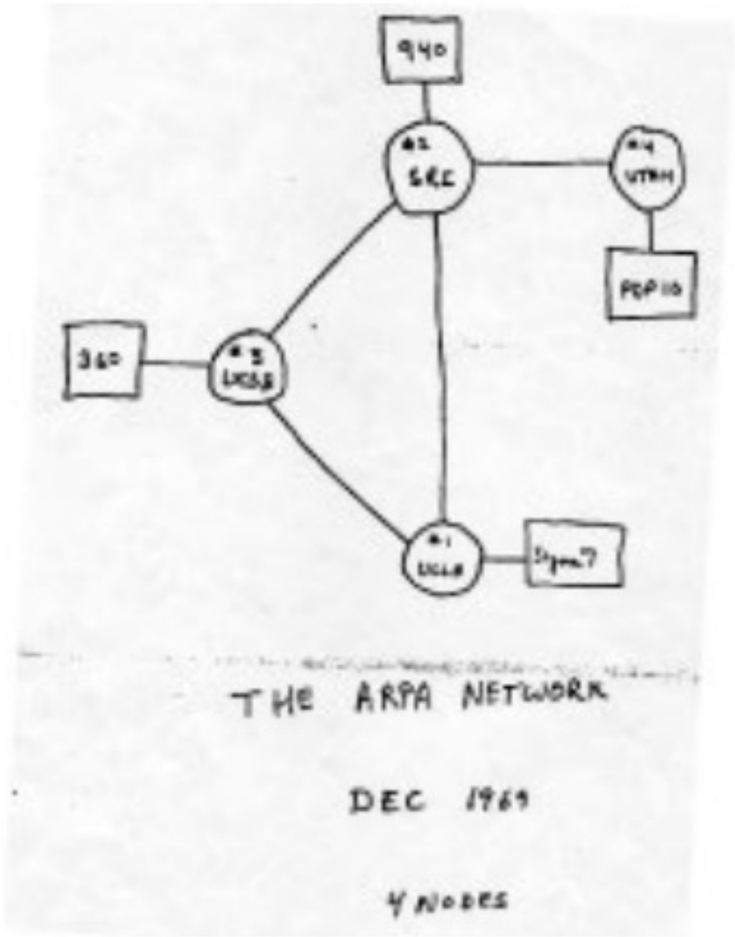

Security through obscurity

Firewalls and Perimeter-Based Trust Models

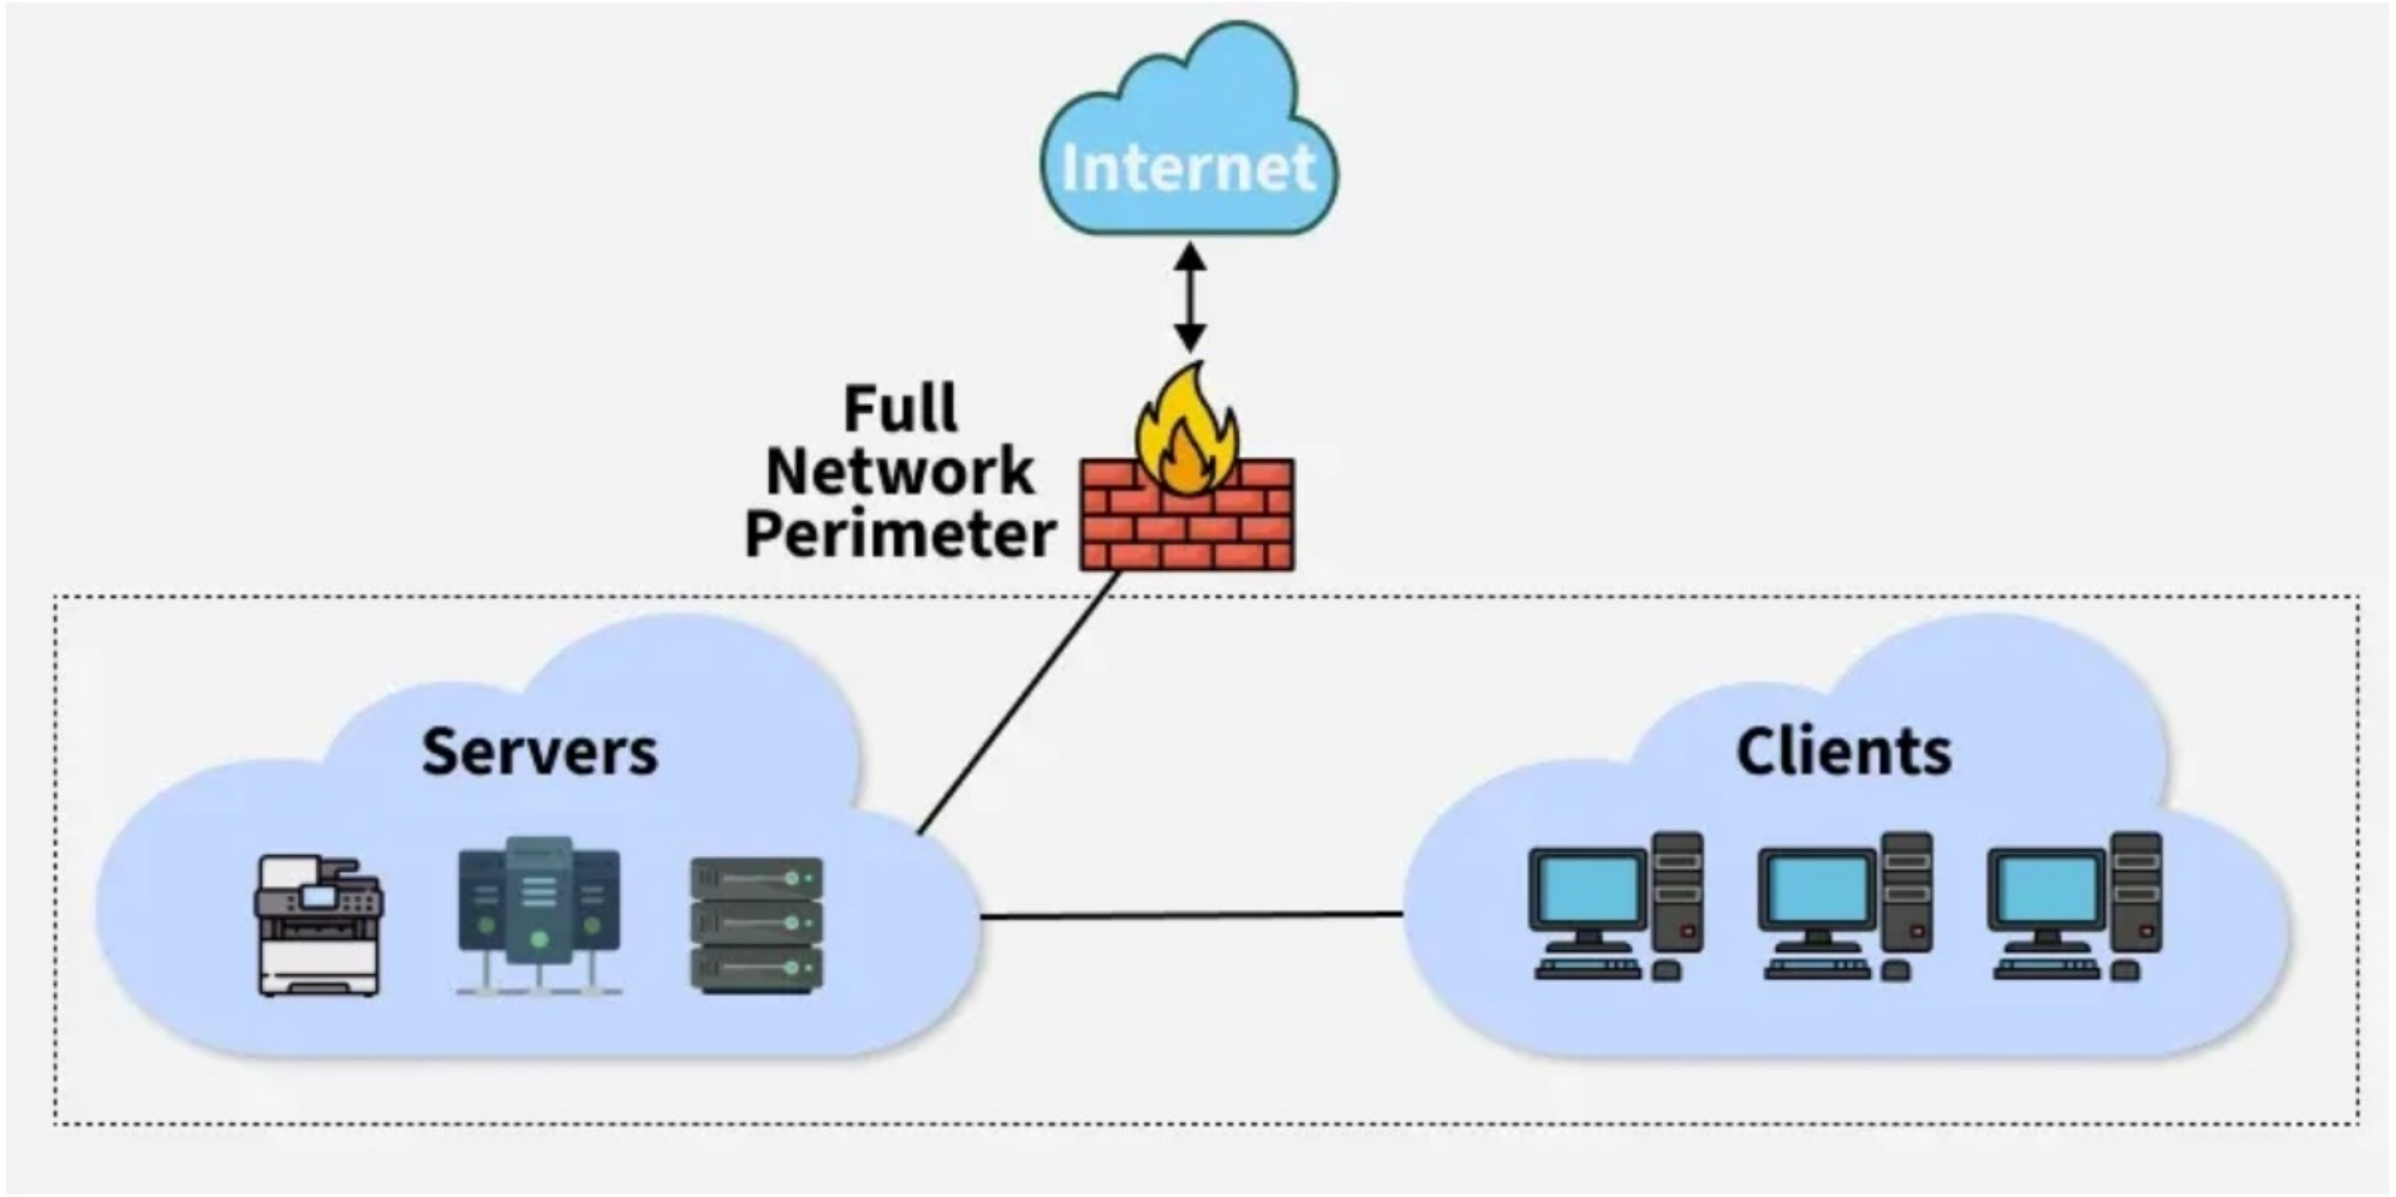

Simple firewall implementation: check for the “evil” bit.

Network Working Group
Request for Comments: 3514
Category: Informational

S. Bellovin
AT&T Labs Research
1 April 2003

The Security Flag in the IPv4 Header

Firewalls [CBR03], packet filters, intrusion detection systems, and the like often have difficulty distinguishing between packets that have malicious intent and those that are merely unusual. The problem is that making such determinations is hard. To solve this problem, we define a security flag, known as the "evil" bit, in the IPv4 [RFC791] header. Benign packets have this bit set to 0; those that are used for an attack will have the bit set to 1.

We talked about another trust model, built on cryptography, earlier tonight.

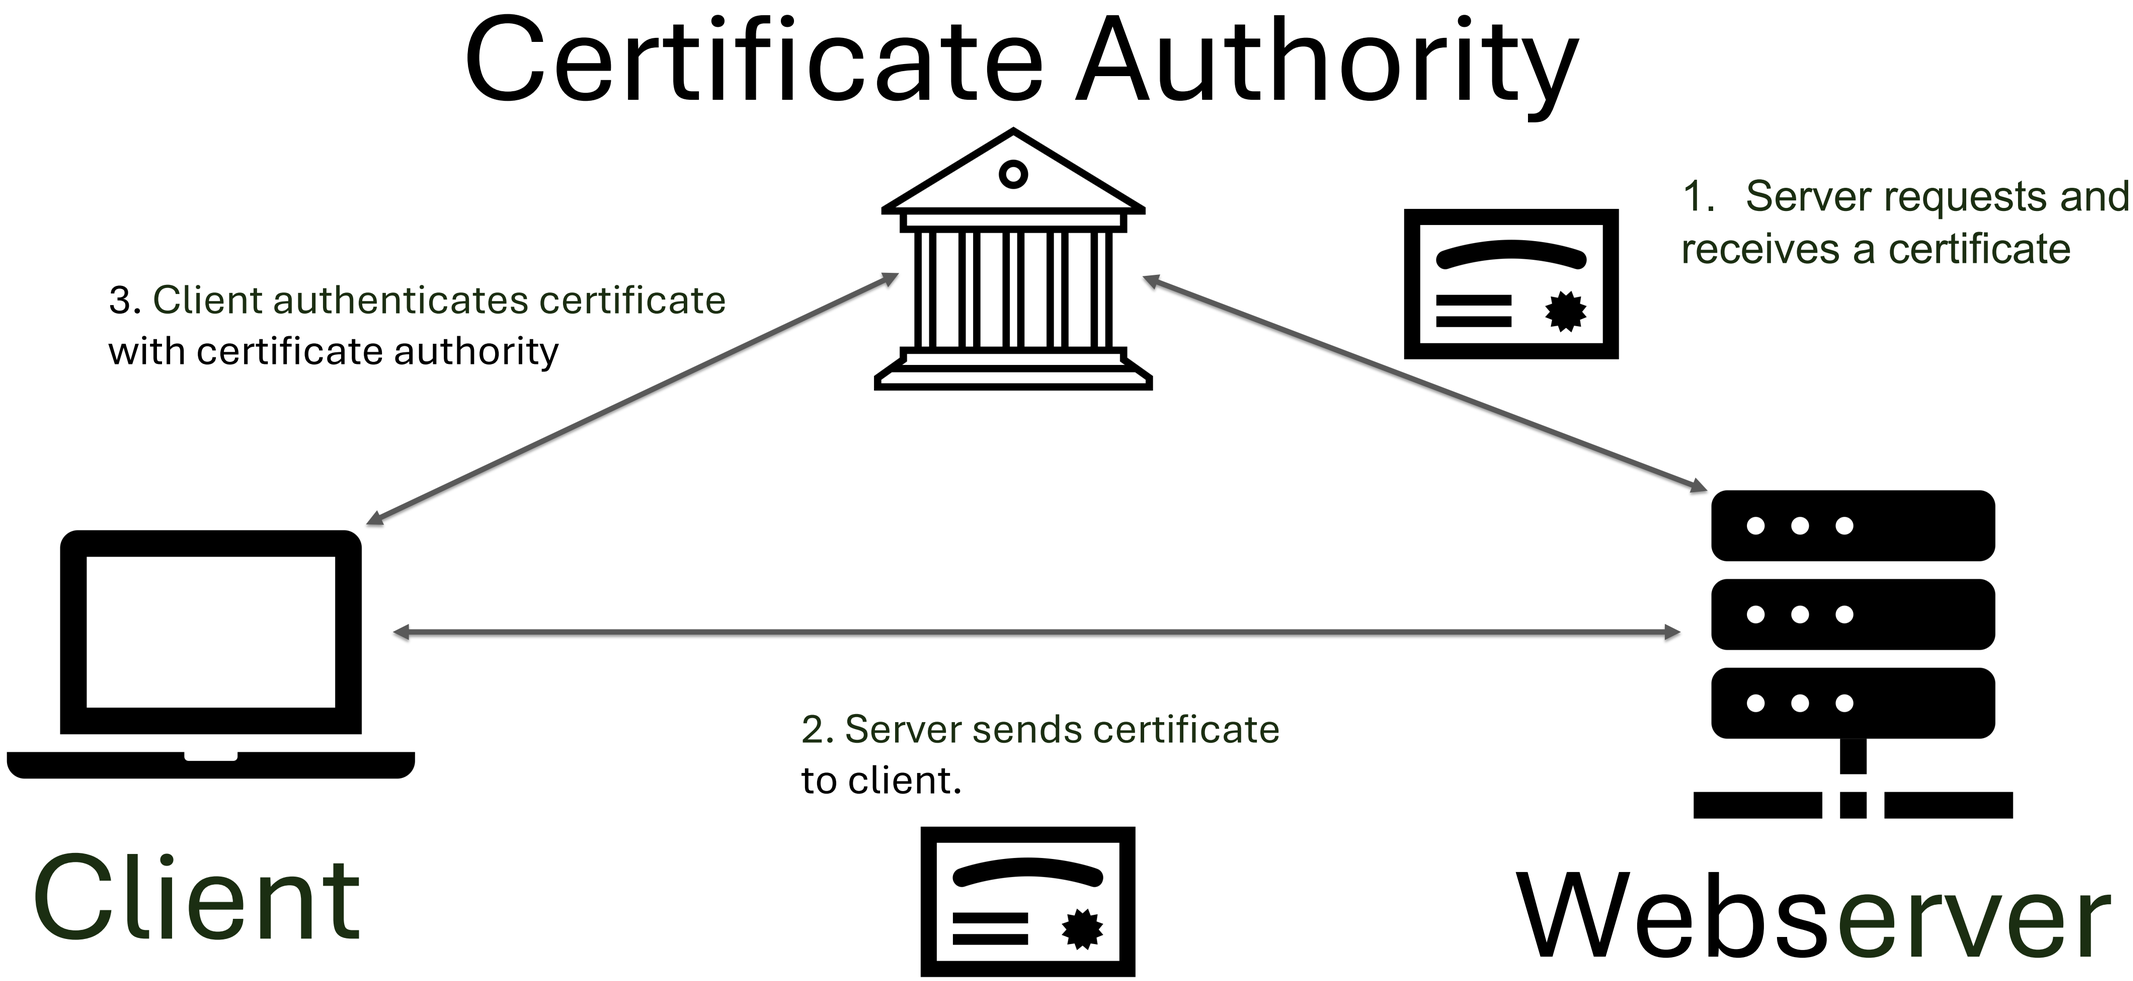

Zero Trust: Check every access from every device to every resource.

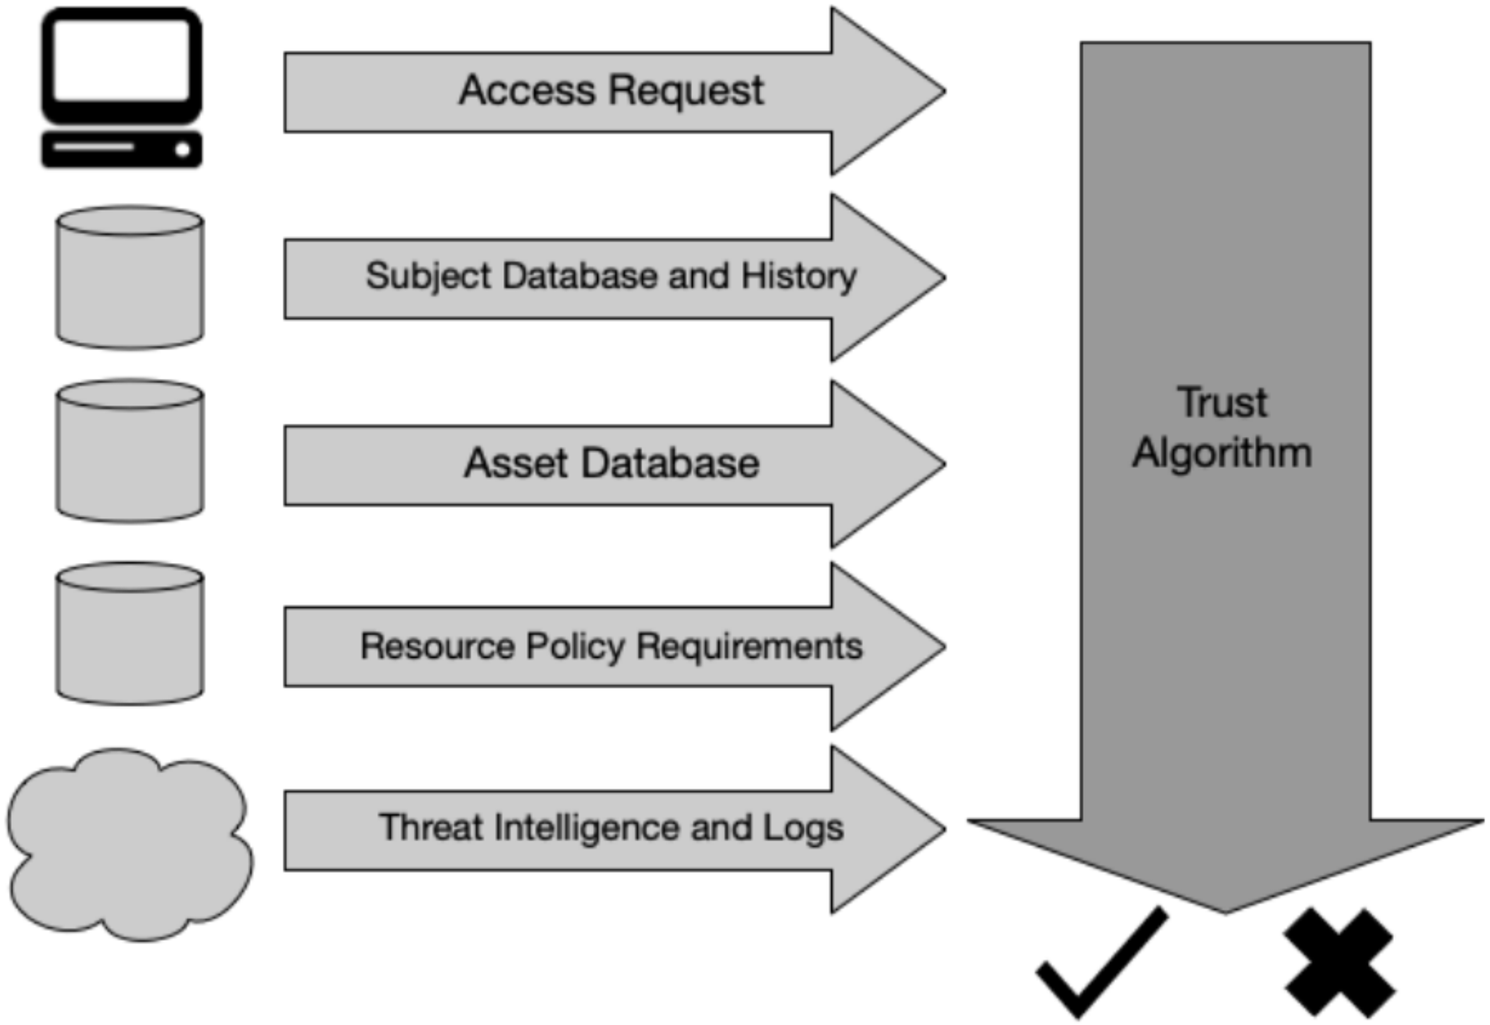

NIST SP 800-207: Zero Trust Architecture

Jan. 30, 2026: NSA releases Zero Trust Guides.

PRESS RELEASE | Jan. 30, 2026

https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/4393480/nsa-releases-phase-one-and-phase-two-of-the-zero-trust-implementation-guidelines/

NSA organizes ZT implementation into seven “pillars.”

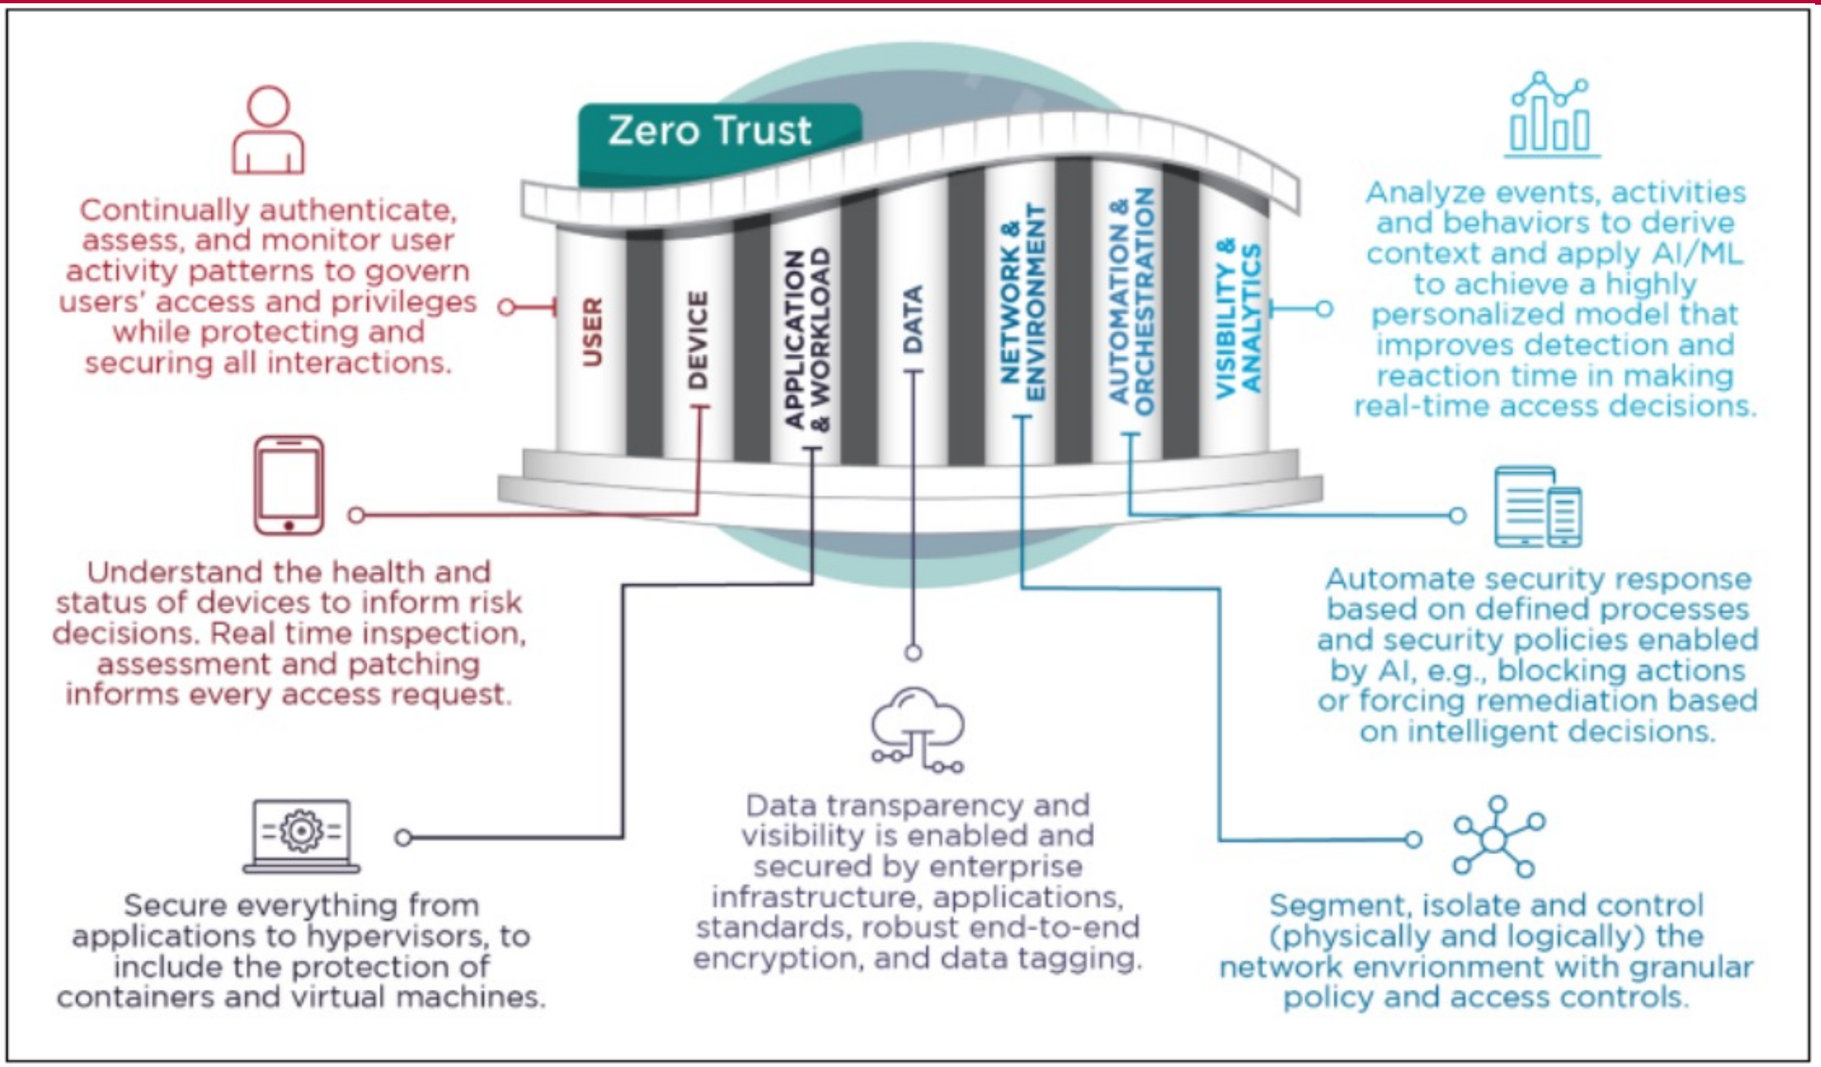

https://media.defense.gov/2026/Jan/30/2003868308/-1/-1/0/CTR_ZIG_PHASE_ONE.PDF

"A Systematic Literature Review on the Implementation and Challenges of Zero Trust Architecture Across Domains"

Systematic Review

A Systematic Literature Review on the Implementation and Challenges of Zero Trust Architecture Across Domains

Sadaf Mushtaq [1], Muhammad Mohsin [2,*] and Muhammad Mujahid Mushtaq [3]

[1] Department of Information Security, National University of Sciences and Technology, Islamabad 44000, Pakistan
[2] Department of Computer Science, Bioengineering, Robotics and Systems Engineering (DIBRIS), University of Genoa, 16146 Genoa, Italy
[3] Department of Electrical and Computer Engineering, University of Victoria, Victoria, BC V8W 2Y2, Canada
* Correspondence: muhammad.mohsin@edu.unige.it

Abstract

The Zero Trust Architecture (ZTA) model has emerged as a foundational cybersecurity paradigm that eliminates implicit trust and enforces continuous verification across users, devices, and networks. This study presents a systematic literature review of 74 peer-reviewed articles published between 2016 and 2025, spanning domains such as cloud computing (24 studies), Internet of Things (11), healthcare (7), enterprise and remote work systems (6), industrial and supply chain networks (5), mobile networks (5), artificial intelligence and machine learning (5), blockchain (4), big data and edge computing (3), and other emerging contexts (4). The analysis shows that authentication, authorization, and access control are the most consistently implemented ZTA components, whereas auditing, orchestration, and environmental perception remain underexplored. Across domains, the main challenges include scalability limitations, insufficient lightweight cryptographic solutions for resource-constrained systems, weak orchestration mechanisms, and limited alignment with regulatory frameworks such as GDPR and HIPAA. Cross-domain comparisons reveal that cloud and enterprise systems demonstrate relatively mature implementations, while IoT, blockchain, and big data deployments face persistent performance and compliance barriers. Overall, the findings highlight both the progress and the gaps in ZTA adoption, underscoring the need for lightweight cryptography, context-aware trust engines, automated orchestration, and regulatory integration. This review provides a roadmap for advancing ZTA research and practice, offering implications for researchers, industry practitioners, and policymakers seeking to enhance cybersecurity resilience.

Academic Editors: Rafael Asorey-Cacheda and Nikolaos Pitropakis

Received: 4 August 2025
Revised: 5 September 2025

Their methodology for lit review:

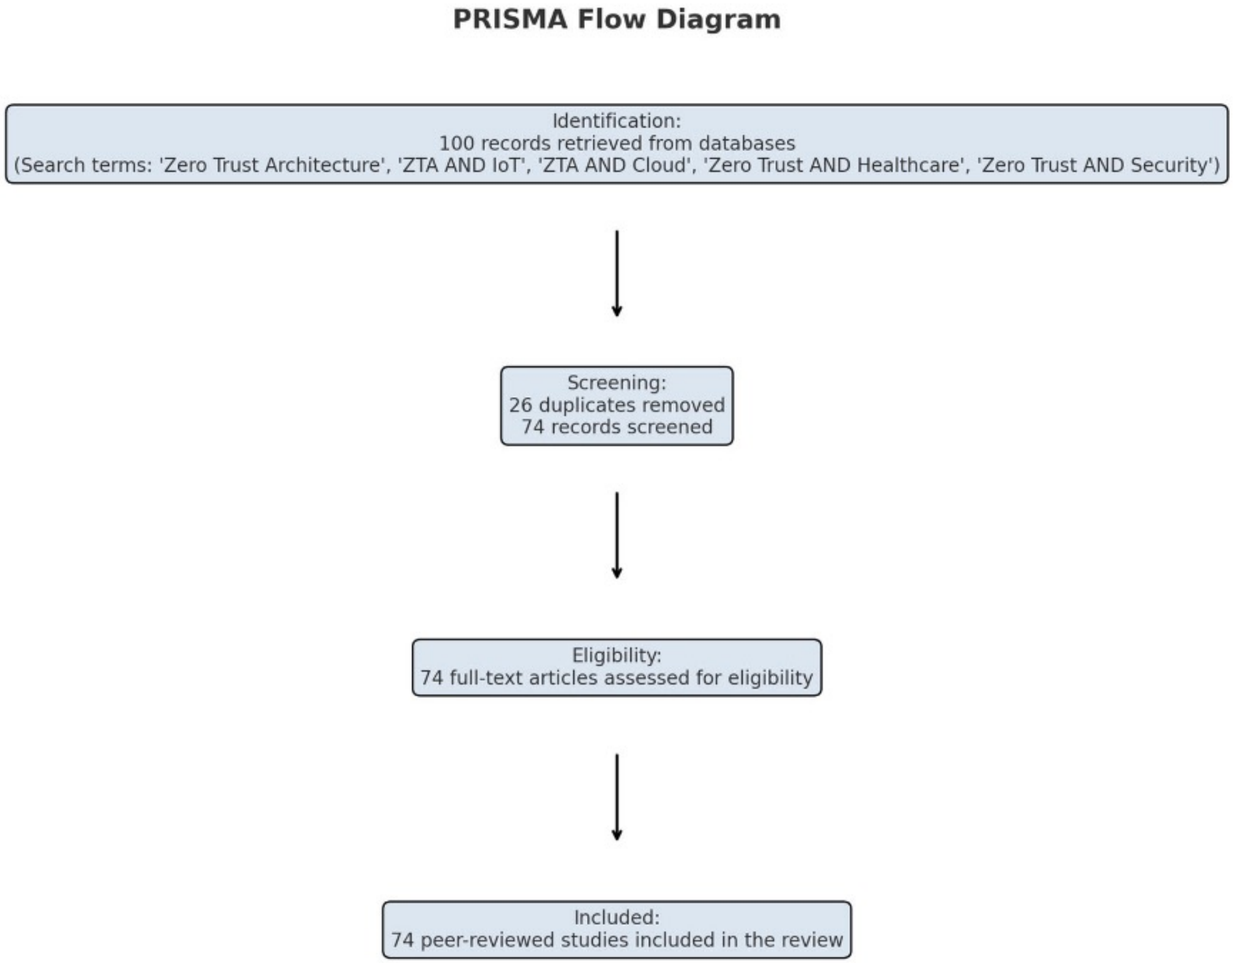

Figure 1. PRISMA flowchart showing the identification, screening, eligibility, and inclusion of studies in this review.

Results of lit review (partial):

Sr No	Year	Paper Id	Applications	Authorization	Authentication	Access Control	Cryptography	Security Gateway	Environmental Perception	Network Segmentation	Audit
5	2024	Zag ElSayed et al. [5]	Healthcare IoT	✓	✓	✓	✗	✓	✓	✓	✓
6	2025	Mohammed A [8]	IoT + Blockchain	✓	✓	✓	✓	✓	✗	✓	✓
7	2024	Nurun Nahar et al. [14]	6G Networks	✓	✓	✓	✗	✗	✗	✓	✓
8	2025	Asif Ali Laghari et al. [15]	IIoT	✓	✓	✓	✗	✗	✗	✓	✓
9	2025	Muhammad Liman Gambo and Ahmad Almulhem [16]	Multiple	✓	✓	✓	✗	✗	✗	✓	✓
10	2025	Shruti Kulkarni et al. [17]	Smart devices	✓	✓	✓	✗	✓	✗	✓	✓
11	2025	Kai Li et al. [9]	IoT + FMs (AI)	✓	✓	✓	✓	✓	✗	✓	✓
12	2025	Izhar Ahmed Khan et al. [6]	SDVs + IoT	✓	✓	✓	✗	✓	✓	✓	✓
13	2025	Lorenzo Gigli et al. [18]	Blockchain + IoT	✓	✓	✓	✗	✗	✓	✓	✓
14	2025	Malek Al-Zewairi et al. [19]	IoT + Traditional	✓	✓	✓	✗	✗	✗	✗	✗
15	2025	Wilson Leite Rebouças Filho [20]	Cybersecurity IAM	✓	✓	✓	✓	✗	✓	✓	✓
16	2025	Xiaokang Zhou et al. [21]	Next-gen network security	✓	✓	✓	✓	✗	✗	✓	✗

Labs

Class Stats as of Sunday, Feb 1

Enrolled in class: 74 students

Completed Lab #0:

Did you purchase the AdaFruit MEMENTO?

Yes!	60 respondents	**98 %**	✓
Not yet.	1 respondent	2 %	

98% answered correctly

Lab #1 — as of Feb 3 at 1:00pm eastern time

Logged into dashboard: 31 students

Registered VM: 24 students

Lab 2 – Create a web server!

Objectives for this lab:

- Install the Nginx webserver on your AWS instance and verify that it works.
- Obtain a Let’s Encrypt server certificate for your web server using the DNS name that you obtained in Lab 1.
- Deploy a simple website with two pages — one available to everybody, one password-protected.

What we will test when you ask for Lab 2 to be graded with the **e11** command:

- We will verify that you are running a web server on port 80.
- We will verify that you are running a “secure” webserver on port 443 with a Let’s Encrypt certificate.
- We will verify that you have properly password-protected the "confidential" web page.

For next week

- Finish Lab #1: Due Feb. 9th
- Start Lab #2: Due Feb. 16th
- Read “Zero Trust Architecture” and comment on Perusall
- Participate in Ed Discussion
- Stay warm!