




Image: NASA Earth Observatory / Lauren Dauphin, Joshua Stevens; USGS Landsat 8

CSINT - BOOK - 2026 - 001

BOOK

OSINT, Step by Step

Open-source investigation through real cases

Free chapters: 1, 2, 3 and 19

Ahmet Göker

First edition, 2026 / CSINT Research

OSINT, Step by Step

Open-source investigation through real cases

Author	Ahmet Göker
Edition	First edition, 2026
Publisher	CSINT Research
Sheets	18 method chapters, 9 real cases
This file	Free chapters: 1, 2, 3 and 19

OSINT, Step by Step

Open-source investigation through real cases

Ahmet Göker

First edition, 2026

CSINT Research, csintresearch.org

CSINT-BOOK-2026-001

This is an independent book. It is not a GIAC or SANS publication and is not endorsed by them. It helps you prepare for the publicly listed exam areas of the GIAC Open Source Intelligence (GOSI) certification, but it contains no exam questions and does not promise a result.

The real cases use public records. Each record is cited where it is used, with its date and its limits. Each image is credited where it appears and listed again at the end of the book. No image in this book was generated by artificial intelligence.

Cover: the Gulf of Suez, 27 March 2021, Landsat 8. NASA identifies the small traces at sea as waiting ships. The image is one of the records examined in Chapter 19. NASA Earth Observatory / Lauren Dauphin, Joshua Stevens; USGS Landsat 8.

The location on the cover is the map point NASA Earth Observatory gives for these images: 29.8256 N, 32.5661 E. The source gives no corner coordinates for the image, so the cover prints none.

© 2026 Ahmet Göker, CSINT Research. All rights reserved. No part of this book may be copied, shared or republished without written permission, except as the law allows. This free file holds the chapters that are free to read and is for personal reading.

Contents

This file holds the chapters of the book that are free to read: 1, 2, 3 and 19. The other chapters are in the full edition, which is listed at the end of this file.

A note from the author	5
How to read this book	7
Before you start	8
GOSI study map	9

Part 1: First, how to think

01	Write the question first	12
02	Separate source and claim	24
03	Prepare before searching	36

Part 6: Real cases

19	Did the queue clear when the Suez Canal reopened in 2021?	48
	Answer key	62
	Images and rights	65
	The full edition	66
	About the author	68

A note from the author

I like puzzles. Laying the pieces out in front of me and working out which corner of the picture is still missing has always done me good. On free evenings I often watch true-crime documentaries, serial-killer cases among them. I want to understand why a person acts the way they do. I also want to see how one small detail can crack a case that went unsolved for years. In the BTK case, a first name and a church name in a floppy disk's file details, searched together, led investigators to one person. In the Gilgo Beach case, years later, the registration of a pickup truck a witness remembered gave investigators a name.

That curiosity brought me to OSINT. Instead of only watching a case being solved, I wanted to look at the same records myself. While I studied for GOSI, I spent eight weeks keeping my own notes. I also built a topic index to use in the exam. This book grew from that work. I wanted to keep the proper technical terms and still explain each idea in small steps. Every chapter has a real record, a diagram and a short exercise, and the last part follows real cases from start to finish.

Some pictures make me angry, and some frighten me. Some claims sound like what I already believe. At such moments, if my feelings or political views come first, I jump to an answer before checking the source properly. I try to catch this while I work.

In my notes I write what I can see, my own reading and the missing information on separate lines. I keep what a photo shows apart from what I read elsewhere. A claim that fits my views gets the same scrutiny as any other.

Sometimes the record I have cannot answer the question. Then I write down what is missing. Readers should be able to see which source my conclusion rests on.

01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

There are no photographs of victims in this book. I do not describe the violence in detail, and I leave out the names of relatives who had nothing to do with the case. In the real crime cases I follow the investigators' steps and the records those steps rest on.

Ahmet Göker


How to read this book

The book has six parts. The first five teach one method per chapter: how to frame a question, keep a record, read a web or email record, and write down what you know. The sixth part applies those methods to real events, from start to finish.

Each method chapter opens with a short summary page. Then come the lessons, a worked example, a point to keep in mind and a short exercise. Try the exercise before you look at the answer; the answers are collected at the back of the book, and each exercise gives the page of its answer.

Technical terms are kept in the form you will meet in tools, records and exams. Each method chapter explains its terms under “Terms in this chapter”, and the glossary lists them all again.

The colours mean the same thing in every diagram. A blue card is what a record shows, a dashed ochre card is your reading of it, which needs support. A violet diamond is a question, a green line means yes and a red line no. The heavy orange card is the result you may write down.



The landscape pages at the end of some chapters are my work screens. I redrew them in the OSINT VM (virtual machine) look of my YouTube videos. The documents, marks, notes and Python output in them are real. I ran the Python code again on my saved files when I made each screen, so you can repeat the numbers from the same files.

In the real cases every record is shown with its date, its source and its limits. When a record cannot answer a question, I say so and stop there.

Every method chapter ends with “In the field: try it yourself”, tasks on real public records that you can repeat today. Under each task is what I saw on the day, a screenshot of the page I opened or the real output of the command I ran. The value I found is marked.

Before you start

Most field tasks need only a browser. “Open” means open the address there; “Search” means enter the query in a search engine. For “Run”, type or paste the command in a terminal and press Enter. “Open + run” needs both. Compare the result with the image under the task. Find text on a page with Ctrl+F, or Command+F on a Mac.

On macOS, Terminal is in the Utilities folder inside Applications. `curl`, `dig`, `whois`, `grep` and `shasum` come with it.

On Windows 11 or supported Windows 10, follow learn.microsoft.com/windows/wsl/install. Open PowerShell as administrator, run `wsl --install`, restart and open Ubuntu. In Ubuntu or Debian, run `sudo apt update` and `sudo apt install curl dnsutils whois jq gnupg`. On Linux, replace `shasum -a 256` with `sha256sum` for the same SHA-256 value.

Chapter 15 needs GnuPG (`gpg`); Chapter 17 uses `jq`. On macOS, see gnupg.org/download and jqlang.org/download, or run `brew install gnupg jq` with Homebrew.

“The same record on the command line” marks an optional command image under a browser task. Chapter 17 explains how to read such commands.

Records can change after 27 September 2026. Write your own date next to each result.

Most tasks use public records. Chapter 3 also tests your browser and optionally checks a link in your own inbox. Do not open or share that link. The tasks do not scan systems or seek another person’s private information.

GOSI study map

I wrote this book for people preparing for GOSI and for everyday research. It is an independent work. It is not an official GIAC or SANS course book. The study map shows which chapters cover the six public exam areas. I wrote the exercises for this book. I cannot know which questions the exam will ask, and I cannot promise a pass.

Exam area	Chapters	What to practise
Method and reporting	1. Write the question first 2. Separate source and claim 18. Write what you know	Write the question, scope, sources and result. Test another possible explanation.
Privacy and OPSEC	3. Prepare before searching 5. Keep what you find 15. Privacy and keys	Explain who can see an action and what data it reveals.
Search and data processing	4. Narrow the search 5. Keep what you find 6. Look at a photo three ways 7. Where? When? 8. Put the rows in order 16. Learn a tool by its job 17. Work with a small file	Build a query. Save a record. Process a file. Check the result with a small sample.
People research	9. Separate account and person 6. Look at a photo three ways 7. Where? When? 13. Which name do email checks verify?	Compare accounts using time and source records. Explain a possible false match.

Exam area	Chapters	What to practise
Company research	10. Find the record behind the brand 14. Whose server? What evidence? 15. Privacy and keys 18. Write what you know	Find the right company. Check a breach claim, a wireless record and business impact separately.
Networks and infrastructure	11. Name, address and DNS 12. Read the name in the certificate 13. Which name do email checks verify? 14. Whose server? What evidence?	Read DNS, IP, registration and certificate data. Separate provider and customer links.

The current public GOSI page lists 75 questions, 2 hours and a 69% pass mark. Check your exam version in your GIAC account. Internet and electronic notes are not allowed in the exam. GIAC also bans printed material that looks like test questions and answers. Do not assume this exercise PDF is allowed in the exam room.

On 16 September 2026 I put the GIAC objectives, the public SEC497 topic list and the 28-page source PDF side by side. This map only links the public areas to chapters. I have not seen GIAC's question bank, and I cannot promise that this book is all you need for the exam.

First, how to think

Before you search, write what you want to learn, separate the source from the claim and prepare your setup.

- 01 Write the question first
- 02 Separate source and claim
- 03 Prepare before searching

01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

Write the question first

Before you search, write one sentence about what you want to find out.

In this chapter

- 1.1** Choose a small question
- 1.2** How do we get an answer?
- 1.3** From data to a decision
- 1.4** In the field: try it yourself

1.1 Choose a small question

In March 2021 the container ship Ever Given ran aground in the Suez Canal. I could have started with “What happened at Suez?” But then I would not have known where to stop. The ship, the canal, trade and insurance all fit inside that question. I split it in two. I wrote the first part like this: “Did the canal reopen on 29 March 2021?”

A record can answer that question. I put the date inside the question too. “It reopened yesterday” points to a different day depending on when you read the report.




Figure 1.1 Splitting the question gives two separate dates. Write each event’s source beside it so the registration date is not taken for the connection date.

1.2 How do we get an answer?

In the Suez work, once I saw that the queue had not cleared on 29 March, I started asking on which day it did clear. This loop, from a question to an answer and on to a new question, is the intelligence cycle. The figure below gives a short definition of each stage, and the table shows what each stage actually meant in my Suez work:

01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

Direction	I wrote two separate questions: did the canal reopen on 29 March 2021, and had the waiting ships cleared that day?
Collection	I gathered NASA's daytime and night satellite images, the Suez Canal Authority's notice of 29 March and ESA's radar comparison.
Processing	I put every record in a table with its publication date and the date it was taken.
Analysis	I compared the records. According to NASA's account, the canal reopened on 29 March. The canal authority's notice reports that the ship was refloated. It adds that navigation would resume once the ship was fully aligned. It does not mention a reopening. By the Leth Agencies figure NASA relays, 184 ships were still waiting on 30 March, so the queue had not cleared on 29 March.
Production	I wrote the result in a short note with its sources and added what I did not know.
Dissemination	I published the note as a research file.
Feedback	If a reader asks "On which day did the queue clear?", the work starts again with that question.




Figure 1.2 If the reader comes back with a response or a new question, the work starts again at Direction.

1.3 **From data to a decision**

When you look at a record, first ask what it is for: the same record can mean different things for different questions. The rows below show how satellite images of the Gulf of Suez lead, step by step, to an assessment.

Data	Small, elongated traces at sea in a satellite image dated 27 March 2021. No interpretation yet.
Information	Comparing the records gives them meaning: the traces are far fewer on 11 March, and NASA identifies them as ships waiting to enter the canal.

Intelligence	You may conclude that while the blockage lasts, more ships will wait and cargo bound for the canal will be late. That is an assessment. Write it with its evidence and its gaps.
Decision	Record who will use the answer, by when, and which actions you are authorised to take. Scope sets the limits of this work.




Figure 1.3 The traces in the 27 March image become ships waiting for the canal once NASA explains them. The question in the dashed frame is not in the record. You are the one asking it.




Figure 1.4 With a narrow, dated question, you know which record to look for.

A worked example

Let's do it together

Your task is to find out whether the Suez Canal reopened on 29 March 2021. You do not need the crew's or the canal staff's personal accounts for that. Start from the notice of the body that runs the canal.

- 01 Write the question with its date: "Did the canal reopen on 29 March 2021?"
- 02 Find the Suez Canal Authority's notice of 29 March. Note its link, its publication date, its headline and its body separately.
- 03 If the headline and the body, or two sources, say different things, write that down plainly. If no record shows that the canal reopened, do not give a definite answer.

The case this example comes from: Chapter 19, Did the queue clear when the Suez Canal reopened in 2021? (page 48).



Figure 1.5 Write the four parts on separate lines. If the headline and the body say different things, put that plainly in your note.

1.4

In the field

Try it yourself

1

Split the question in two

Open <https://www.iana.org/domains/root/db/tr.html>

Do Split “When did Türkiye get the internet?” into two: “When did .tr enter the IANA register?” and “On which day did the first internet line open?” Scroll to the very bottom of this page from IANA. IANA is the body that keeps the master list of domain extensions such as .com and .tr. Note the “Registration date” and “Record last updated” lines. Then find the day of the first internet connection in TÜBİTAK’s article at bilimgenc.tubitak.gov.tr/makale/turkiyede-internet-25-yasinda and compare it with the “Registration date”.



The 1990-09-17 registration date, boxed in red at the bottom of IANA’s .tr page. IANA, screenshot 27.09.2026

What I found The bottom of the page read “Registration date 1990-09-17” and “Record last updated 2026-07-05”. TÜBİTAK’s article of 12 April 2018 dates the first link to 12 April 1993. It was a 64 Kbps leased line between METU and NSFNet in the US. The same article also mentions BITNET-EARN access in 1987 and an X.25 test link in September 1992. By that count .tr entered the register 2 years, 6 months and 26 days before the first line opened.

Watch out “Registration date” is the day the extension entered the register. No connection had to exist that day. “Record last updated” is the date of the latest change.

2

Write down the figure's date and limits

Open <https://veriportali.tuik.gov.tr/tr/press/53899>

Do Instead of “What is Türkiye’s population?”, write “How many people were resident in Türkiye on 31 December 2025?” Note the “Yayın Tarihi” (release date) line at the top of the bulletin. Separately, note the date in the first paragraph that the figure refers to. Read note (1) near the bottom of the page.



The TÜİK bulletin on my phone. The line boxed in red is the release date, 9 February 2026, 10:00. Türkiye İstatistik Kurumu (TÜİK), screenshot 27.09.2026

What I found The top of the bulletin read “Yayın Tarihi: 09 Şubat 2026 10:00” (released 9 February 2026, 10:00). The first paragraph gives 86,092,168 residents as of 31 December 2025. Note (1) says that Syrians under temporary protection are not counted. Foreigners are not counted either if they stay on visas or residence permits of under 90 days for reasons such as tourism.


Watch out The figure does not answer “How many people are in Türkiye today?”. It is a year-end record and it leaves some groups out.

3

Find the week of the country's name change in the archive

Open <https://web.archive.org/web/20220531130855/https://www.un.org/en/about-us/member-states>

Do Write the question with its date: “How did the UN member states page spell Türkiye’s name on 31 May 2022?” Search the page (Ctrl+F) for Tunisia and read the line just below it. Then change 20220531130855 (the moment the archive took this copy, 31 May 2022 13:08:55) in the address to 20220607132805 and look at the same line.



```
Terminal: Wayback Machine (CDX) 27.09.2026

$ curl -s 'https://web.archive.org/cdx/search/cdx?url=www.un.org/en/about-us/member-states&from=20220531&to=20220607&fl=timestamp,statuscode'
20220531055112 200
20220531055631 200
20220531104115 200
20220531130855 200
20220607132805 200
$ for t in 20220531130855 20220607132805; do echo $t; curl -s https://web.archive.org/web/${t}/https://www.un.org/en/about-us/member-states | grep -A 2 -E '>(Turkey|Türkiye)<' | sed -e 's/<[^>]*>/g' -e 's/^[:space:]*//'; | grep .; done
20220531130855
Turkey
Date of Admission: 24-10-1945
20220607132805
Türkiye
Date of Admission: 24-10-1945
```

The same record on the command line

The CDX query of 27 September 2026 and the two captures. I highlighted the Turkey and Türkiye lines with both timestamps.

What I found The capture of 31 May 2022, 13:08 UTC read “Turkey Date of Admission: 24-10-1945”. UTC is universal time, three hours behind Türkiye. The capture of 7 June 2022, 13:28 UTC read “Türkiye Date of Admission: 24-10-1945”. The archive holds no capture between the two. A UN in Türkiye story dated 3 June 2022 says the name-change letter was received on 1 June. The change took effect from that moment.

01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

Watch out

The archive cannot give you the moment the page changed. You only have the gap between two of its captures, here seven days.

Keep in mind

While you search, go back to the question you wrote at the start and check whether what you found answers it.



Your turn

How would you narrow the task “Collect information about the Suez Canal”?

Answer: page 62



Puzzle

symbolics.com is known as the first .com name ever registered. Its RDAP record (its registration data, at rdap.verisign.com/com/v1/domain/symbolics.com) gives a registration time of 1985-03-15T05:00:00Z. The expiry time is 2027-03-16T04:00:00Z. The T separates the date from the time, and the Z at the end means the time is in UTC. Why does the same record say 05:00 in one place and 04:00 in the other? Hint: think about daylight saving time.

Answer: page 62




Figure 1.6 Add one part at each step. In the end the question can be checked in a single record, such as NASA's satellite image of that day.



Check what you learned

The task asks whether a service outage affects a company. Your first list has 500 staff names. What do you do?

Answer: page 62

Terms in this chapter

OSINT	Open source intelligence: answering a question with evidence drawn from publicly available sources.
Intelligence requirement	The question the research has to answer.
Scope	The limits of the work: topic, dates and permitted actions.
Seed / Pivot	The clue you start from / using that clue to move to a



new record.

Stop condition

A rule, set at the start, for when to stop searching.

Data / Information / Intelligence

Raw records / records given context and meaning / an assessment that supports a decision.

Sources

- 1. GIAC, GOSI objectives
<https://www.giac.org/certifications/open-source-intelligence-gosi>
- 2. NASA Earth Observatory, Traffic Jam on the Suez Canal (31 March 2021)
<https://science.nasa.gov/earth/earth-observatory/traffic-jam-on-the-suez-canal-148114/>
- 3. Suez Canal Authority, Successful Refloating of EVER GIVEN (29 March 2021)
https://www.suezcanal.gov.eg/English/MediaCenter/News/Pages/nav_29-03-2021.aspx

Separate source and claim

Even a trusted source can be wrong. NASA's article of 31 March says the canal reopened, but NASA did not observe that itself. Its reference list names the canal authority's notice and Reuters. Trace each sentence to the record behind it.

In this chapter

- 2.1** Keep three separate lines
- 2.2** Find the source behind the copies
- 2.3** Assess the source and the claim separately
- 2.4** Test your first idea
- 2.5** In the field: try it yourself

2.1 Keep three separate lines

A library notice reads “Closed on Sunday”. In your notes, keep what you saw, what you infer from it and what you are taking for granted on separate lines:

Fact	The notice I saved reads “Closed on Sunday”.
Inference	The library may be closed this Sunday too.
Assumption	I accept that the notice still applies.

Only the first line rests on a record. The other two do not yet confirm this Sunday’s hours, so check when the notice was published.

2.2 Find the source behind the copies

Five accounts may share a photo of the same notice. If they all copy one notice, you have one source, repeated five times. Follow the copies back to where the notice first appeared. On the way you may find records that do not fit your idea. Note those too. Looking only for information that supports the answer you expect is called confirmation bias.




Figure 2.1 Follow the copies back to the first source. If a record turns up that does not fit your idea, note it as well.

2.3 Assess the source and the claim separately

Here you ask two separate questions. Source reliability looks at the source’s track record and how it got the information. Information credibility asks what this one claim rests on. Keep them apart, because a reliable source can still pass on something wrong.

Admiralty / NATO-style grading	Sources are graded with letters from A to E, information with numbers from 1 to 5. F and 6 mean that the source or the information could not be judged. Organisations may define the scale differently, so check the definitions where you work.
F6	There is not enough information to judge either the source or the claim. The claim may still be true.
Confidence	How strongly the evidence supports the conclusion. When you write low, moderate or high, add the reason.
Impact	How much harm or change would follow if the claim were true? Write this separately from your confidence: a large impact does not make a claim more likely.





Figure 2.2 Grade the source and the claim separately. Organisations may define the scale differently, so check the definitions where you work.

2.4 Test your first idea

Once you have written an explanation, look for a record that could weaken it. Two very similar sites might have one owner, or they might simply use the same design agency. Ask which finding would tell these explanations apart. The biases below can tempt you to skip this test.

Confirmation bias	Selecting only information that supports your first idea.
Anchoring / Premature closure	Giving the first clue too much weight / deciding too early and stopping.
Availability / Authority bias	Giving recent or vivid events too much weight / trusting a famous source without checking.
Mirror imaging	Assuming another person thinks like you.
Correlation / Causation	Two events appearing together / one event causing the other. Seeing them together leaves the cause open. You look for the cause in a separate record.



Look for a finding that would tell these two apart.

Figure 2.3 Both rows hold the same likeness. However long you look at it, you cannot pick one of the two explanations.



Figure 2.4 Start from the claim and take each branch in turn. Answer the question in each box, and base your confidence only on the records you find.

A worked example

Let's do it together

For Suez I had three records of the same event. The first was the Suez Canal Authority's notice of 29 March 2021. The second was the Sentinel-1 radar comparison that the European Space Agency (ESA) published on 26 March. The third was NASA's article of 31 March. They do not show the same thing. The notice is the canal authority's own statement.

ESA and NASA made two separate observations with different satellites. If a news site republishes NASA's image, all I gain is a copy of the same record.

- 01** Write down where each record comes from. The notice is the canal authority's own statement. The radar comparison is ESA's. The day and night images are NASA Earth Observatory's (made from USGS Landsat 8 and Suomi NPP VIIRS data).
- 02** Separate each publication date from the date of the observation behind it. ESA published on 26 March, but it compares images from 21 and 25 March. NASA published on 31 March, and its images were

taken between 1 February and 29 March.

- 03 Keep independent observations apart from copies. NASA and ESA looked with different satellites, so their images support each other. NASA's sentence about the reopening does not come from its own observation, though. Its reference list names the canal authority's notice and Reuters.




Figure 2.5 Each record has two dates. Find the observation date first, then write the publication date on its own line.

From a real record

Two pages, the same map

In the Beirut study, I compared pages from NASA and NASA Earth Observatory and looked at which measurements each one relied on.

- The map comes from analysis by ARIA and the Earth Observatory of Singapore.
- The same analysis also appears on a different background, so the source chain leads back to the same work.




Figure 2.6 Read the question and follow the matching branch.

What I can write

I keep both links in my notes. I do not count them as two independent pieces of evidence.

The full case: Chapter 21, in the full edition.

Where I stop

The second page gave me no new observation.

2.5

In the field

Try it yourself

1

Find the earliest record of a quote

Open <https://quoteinvestigator.com/2017/03/23/same/>

Do Open the page on the saying “Insanity is doing the same thing over and over again and expecting different results”, usually credited to Einstein. Write down its earliest known record. The page calls it the “earliest strong match”. Also write down the first year the page says it was tied to Einstein. Also read the “Update History” note near the end of the page.

```

Terminal: curl + grep 27.09.2026

$ curl -sL --compressed https://quoteinvestigator.com/2017/03/23/same/
| sed -e 's/<[>]*>/g' | grep -oE "There is no substantive evidence
[^.]*\.|The (second)?earliest strong match known to QI appeared in
[^.]*|By 1990 the saying was being attributed to Einstein\.|Update Hi
story: On July 31, 2019 [^.]*\."
There is no substantive evidence that Einstein wrote or spoke the stat
ement above.
The earliest strong match known to QI appeared in October 1981 within
a Knoxville, Tennessee newspaper article describing a meeting of Al-An
on, an organization designed to help the families of alcoholics
The second earliest strong match known to QI appeared in a pamphlet pr
inted by the Narcotics Anonymous organization in November 1981
By 1990 the saying was being attributed to Einstein.
Update History: On July 31, 2019 the October 11, 1981 citation was add
ed to the article.
$ curl -sL https://www.nobelprize.org/prizes/physics/1921/einstein/bio
graphical/ | grep -oE "He died on [^.]*\."
He died on April 18, 1955 at Princeton, New Jersey.

```

The same record on the command line

The highlighted lines hold the earliest record, the sentence tying the saying to Einstein by 1990 and the update note. I pulled them with curl on 27 September 2026.

What I found The page says there is no “substantive evidence” that Einstein wrote or said it. The earliest strong match it knows of is a Knoxville News-Sentinel report of 11 October 1981. It covers an Al-Anon meeting, where an unnamed attendee says it. The saying was tied to Einstein by 1990, in the Austin American-Statesman. According to the Nobel Prize site, Einstein had died on 18 April 1955.


Watch out *The “earliest known” record need not be the first time it was said. The page notes that the 11 October 1981 citation was added on 31 July 2019. Until then the earliest strong match was a Narcotics Anonymous pamphlet of November 1981.*

2

Find the edit behind the copies

Open https://en.wikipedia.org/w/index.php?title=Maurice_Jarre&action=history&offset=20090401000000&limit=60

Do Open this history page of composer Maurice Jarre's Wikipedia article. Find the line with "02:29, 30 March 2009". At the start of that line, open the "prev" link (compare with the previous version). In the comparison that opens, read the "Quotes" section the edit adds. Then read the correction note at the bottom of the Guardian obituary of 31 March 2009: theguardian.com/film/2009/mar/31/maurice-jarre-obituary



```
Terminal: Wikipedia API (compare) 27.09.2026

$ curl -s 'https://en.wikipedia.org/w/api.php?action=compare&fromrev=280558491&torelative=prev&prop=ids|timestamp|size&format=json&formatversion=2' | jq -c '.compare | {torevid, totimestamp, fromsize, tosize}'
{"torevid":280558491,"totimestamp":"2009-03-30T02:29:43Z","fromsize":11583,"tosize":12318}

$ curl -s 'https://en.wikipedia.org/w/api.php?action=compare&fromrev=280558491&torelative=prev&prop=diff&format=json&formatversion=2' | jq -r '.compare.body' | grep 'diff-addedline' | grep -v '<br />' | sed -e 's/<[^>]*> //' -e 's/^ *//'
==Quotes==
[...]
One could say my life itself has been one long soundtrack. Music was my life, music brought me to life, and music is how I will be remembered long after I leave this life. When I die there will be a final waltz playing in my head and that only I can hear.
[...]
```

The same record on the command line

The 02:29 UTC edit of 30 March 2009 and the "Quotes" section it added, highlighted in the comparison I pulled from the Wikipedia API on 27 September 2026.

What I found The history shows an edit at 02:29 UTC on 30 March 2009. The edit added a "Quotes" section, with a line that runs "my life itself has been one long soundtrack". The Guardian obituary is dated 31 March 2009. Its note of 3 April 2009 says the piece opened and closed with that invented quote. According to the note, the quote appears to have been deliberately inserted into Jarre's Wikipedia entry and copied from there to various sites.


Watch out A quote printed in several papers is not confirmed by that. If they all rest on the same unsourced edit, you have one source.

3

Separate the observation date from the publication date

Open <https://science.nasa.gov/photojournal/pale-blue-dot-revisited/>

Do Note the “Image Addition Date” line and the sentence that gives the moment the frames were taken (it begins “They were taken at”), separately. Then open the PIA00452 link in the text (the page of the original 1990 version) and read the same line.



The same record on the command line

Lines I pulled with curl from two NASA pages on 27 September 2026. I marked the 2020 addition date and the sentence dating the frames to 14 February 1990 in yellow.

What I found The 2020 version’s page read “Image Addition Date: 02/12/2020”. Its text says the frames were taken at 4:48 GMT on 14 February 1990, 34 minutes before Voyager 1 switched off its cameras for good. The original, PIA00452, reads “Image Addition Date: 09/12/1996”. The 2020 image is not a new observation. It is the same 1990 data reprocessed with modern software, with the brightness of the colour channels rebalanced.

Watch out

“Image Addition Date” is the day the file entered the catalogue. The day the picture was taken is not in that line. This page writes dates month/day/year, so 02/12/2020 is 12 February 2020.

Keep in mind

For every copy, check which record it rests on.

☐ **Your turn**

Does F6 in the Admiralty system mean “false”?

Answer: page 62

☐ **Puzzle**

In a column dated 4 May 2009, the Guardian’s readers’ editor wrote about the invented quote. According to the column, the second time it was added to Wikipedia it stayed for only six minutes. The third time it stayed for about 25 hours. According to the Maurice Jarre page history for 30 and 31 March 2009, which insertion stayed for six minutes? Hint: on the history page from task 2, look at the plus or minus byte count on each line. Then use “prev” to check the lines that add or remove the quote.

Answer: page 62

☐ **Check what you learned**

Three articles use the same anonymous post. How do you count the sources?

Answer: page 62

Terms in this chapter

Fact	Information that a record you hold directly supports.
Inference	A conclusion you draw from the evidence.
Assumption	Something you accept as true for now.
Confidence	How strongly the evidence supports the answer.
Circular reporting	Many reports repeating one source, so that copies look like independent confirmation.
Source reliability / Information credibility	How far the source can be trusted / how well one claim is supported.
Admiralty / F6	A system that grades the source with a letter and the information with a number. F6 means neither can be judged.
Bias / Competing hypotheses	A habit that pulls your thinking one way / different explanations of the same evidence, weighed against each other.

Sources

1. NASA Earth Observatory, Traffic Jam on the Suez Canal (31 March 2021)
<https://science.nasa.gov/earth/earth-observatory/traffic-jam-on-the-suez-canal-148114/>
2. Suez Canal Authority, Successful Refloating of EVER GIVEN (29 March 2021)
https://www.suezcanal.gov.eg/English/MediaCenter/News/Pages/nav_29-03-2021.aspx
3. ESA, Suez Canal traffic jam seen from space (26 March 2021)
https://www.esa.int/ESA_Multimedia/Images/2021/03/Suez_Canal_traffic_jam_seen_from_space

Prepare before searching


The moment you open a page, the site may record your visit, so prepare your working setup before you search.

In this chapter

- 3.1** Know the tool's limit
- 3.2** Does the action reach the target?
- 3.3** Check what you expose before you start
- 3.4** Before opening a file
- 3.5** In the field: try it yourself

3.1 Know the tool's limit

A separate browser profile keeps cookies apart from your personal profile, but it does not hide your IP address. A VM (a virtual computer running inside your own) gives you a separate working environment, but it does not make you anonymous. A VPN changes the exit IP address that sites see. If you sign in to a personal account, the site still sees that account. Incognito only limits the history kept on your device. The site still sees you.



Before you start, check which account you are signed in to and which connection you use.

Figure 3.1 Each row shows what one tool changes and what it leaves exposed. Even when you combine them, check your account and connection separately.

3.2 Does the action reach the target?

A data provider may already have stored a search result. Reading that stored result is not the same as opening the live site. When you read it, your request goes to that provider. When you open the live site, your

request goes to the target’s server. The line is not always clear, though. A screenshot tool visits the site to take its picture.

An archived page may load an image from the live site. Third parties can log your visit too. So look past a tool’s label, however “passive” it sounds, and check what request it actually sends and where.

3.3 **Check what you expose before you start**

A separate research account still leaves traces. Your IP address, open sessions, browser features and the details inside files can tie the research back to you or your organisation. Before you start, write down what data goes to whom.

Browser fingerprint / Cookies	Browser features such as language and screen size can combine into a trace that sets you apart. Cookies can tie a new visit to your earlier sessions.
Sock puppet	A separate account used only for research. It must not impersonate a real person. Keep your personal sessions, photos and account recovery details (backup email, phone number) out of it.
VM / VPN / Incognito	Each changes one thing. A VM separates the work environment. A VPN changes the exit IP address. Incognito limits the history kept on your device. None removes every identity clue.
Pre-flight check	A short check before you start: permission, IP address, DNS (the system that turns domain names into IP addresses), open sessions, browser extensions, autofill and the folder where evidence will be saved.

3.4 **Before opening a file**

Do not trust the extension of an unexpected file: a name ending in “.pdf” can hide another type of file. Keep the original unchanged. Take its hash (a short value calculated from the file’s content, which changes if the file changes). Then work on a copy. First examine the copy without running its code.

This is called static analysis. If you open it in an isolated environment, check that environment’s network access, its clipboard and any folders it shares with your own computer.

Tracking pixel / Canary token	An image loaded from a remote server, or a marked file, can send a request the moment it is opened. That request can reveal when it was opened and from which IP address.
Public upload	A file sent to a public scanning service may be shared with others, so do not upload sensitive evidence without approval.
Retention / Need to know	Decide at the start how long you will keep the data and who needs to see it. Only people who need to know should have access.




Figure 3.2 The original is kept and the setup is limited before any of the file’s code runs. Some files send a request the moment they open, so check network access separately.




Figure 3.3 In each row, the left box is what you do and the right box is who sees the request. The line is not always clear. An archived page may load images from the live site, and a screenshot tool visits the site. If you need contact with the target, check permission and scope first.

A worked example

Let's do it together

In Chapter 22 I look at the domain that WannaCry checks. Every request I send to that domain shows the server behind it my IP address and the time of the request. Before I look at an address linked to malware, I plan what I give away and where each request goes.

- 01** Check your permission and the scope of the work. Keep your personal accounts and browser sessions out of it.
- 02** Start with records that never reach the target. You read the RDAP record from the registry server and the certificates from the CT logs. So these requests do not go to the domain's own server.
- 03** If you need a request to the live server, ask only for the headers (an HTTP HEAD request) and do not download the page. Send it from a separate work environment and note the time in UTC.




Figure 3.4 Start with records that never reach the target. The one request to the live server comes last, and it shows the other side your IP address and the time.

3.5

In the field

Try it yourself

1

Look at your browser's fingerprint

Open <https://coveryourtracks.eff.org/>

Do Open the page in the browser you use for research. “Test Your Browser” sends browser characteristics to EFF; if you do not want that, skip the test and read my example result below. If you run it, you can turn off the real tracking company option. Note whether the result says “unique”. Note its bit value too (a higher number means the browser is easier to distinguish). Then note the Time Zone and Language lines.

01

02

03

04

05

06

07

08

09

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

https://coveryourtracks.eff.org/results?&aat=1&a=1111&t=1111&dnt=1111&fpi_whorls=...27.09.2026

Here are your Cover Your Tracks results. They include an overview of how visible you are to trackers, with an index (and glossary) of all the metrics we measure below.

Our tests indicate that you are not protected against tracking on the Web.

IS YOUR BROWSER:

Blocking tracking ads?	No
Blocking invisible trackers?	No
Protecting you from fingerprinting?	Your browser has a unique fingerprint

Still wondering how fingerprinting works?

LEARN MORE

Note: because tracking techniques are complex, subtle, and constantly evolving, Cover Your Tracks does not measure all forms of tracking and protection.

Your Results

Your browser fingerprint **appears to be unique among the 348,199 tested in the past 45 days**.

Currently, we estimate that your browser has a fingerprint that conveys **at least 18.41 bits of identifying information**.

“unique”, the 348,199 browsers and the 18.41 bits, boxed in red in the Cover Your Tracks result. EFF Cover Your Tracks (CC BY), screenshot 27.09.2026

What I found The test result was “Your browser has a unique fingerprint”. The browser I used was unique among the 348,199 browsers tested in the past 45 days and carried at least 18.41 bits of identifying information. The details showed the time zone “Europe/Istanbul” and the language “tr-TR”. According to the page, only one browser in 221.36 has “tr-TR” as its language value.

Watch out Your name is not in the result. But with a VPN on, the time zone and language stay the same. The browser sends those two values, and the VPN does not touch them.

2

Read the names inside a PDF

Open https://commons.wikimedia.org/wiki/File:JFKennedy_on_Ataturk_1963.pdf

Do Scroll to the “Metadata” table (the details stored inside the file, apart from its visible text) at the bottom of the page. Write down the “Short title”, “Author”, “Software used” and

“Conversion program” rows. Then look at the upload date under “File history”.

Terminal: Wikimedia Commons API27.09.2026

```
$ curl -s 'https://commons.wikimedia.org/w/api.php?action=query&titles=File:JFKennedy_on_Ataturk_1963.pdf&prop=imageinfo&iiprop=timestamp|metadata&format=json&formatversion=2' | jq -r '.query.pages[0].imageinfo[0] | "timestamp: \(.timestamp)", (.metadata[] | select(.name | IN("Title", "Author", "Creator", "Producer", "Page size")) | "\(.name): \(.value)")'
```

timestamp: 2007-10-23T10:24:36Z
Title: JFKennedy1963_text.PDF
Author: mum
Creator: JFKennedy1963_text - Microsoft Word
Producer: Acrobat PDFWriter 4.05 for Windows NT
Page size: 596 x 842 pts (A4)

The same record on the command line

The Author and Creator rows are highlighted. The page's table calls Creator “Software used”. I pulled the metadata from the Commons API on 27 September 2026.

What I found The table read Short title “JFKennedy1963_text.PDF”, Author “mum” and Software used “JFKennedy1963_text - Microsoft Word”. The Conversion program row read “Acrobat PDFWriter 4.05 for Windows NT” and the Page size row “596 x 842 pts (A4)”. The file was uploaded to Commons on 23 October 2007. The text of Kennedy’s 1963 message was converted from a Word document to PDF with Acrobat PDFWriter. The file still carries a file name and a user name.


Watch out The “Author” field shows the user name recorded by the software that made the file. It does not show the author of the text, so “mum” here is not Kennedy.

3

Spot the ID in a newsletter link

- Open
- https://mailchimp.com/developer/marketing/docs/e-commerce/
- Do
- Search the page (Ctrl+F, or Command+F on Mac) for mc_eid and read its paragraph. If you have a newsletter in your inbox, right-click a link without opening it, or long-press it on a phone. Copy its address and note the parameters after “?”

(extra information sent to the server, separated by &). If you have no newsletter, use the mc_eid example on the page. Do not share your own link address.



```
Terminal: curl + grep 27.09.2026

$ curl -sL https://mailchimp.com/developer/marketing/docs/e-commerce/
| sed -e 's/<[^>]*>//g' | grep -oE 'The mc_eid parameter is [^]*\.'|To
find a contact[^]*\.'|For example, if your subscriber[^]*\?unique_ema
il_id=[0-9a-f]*|The mc_cid parameter is [^]*\.'
```

The mc_eid parameter is a unique Mailchimp ID that identifies the recipient's email address, which you can use to associate orders with a contact.

To find a contact's email address, make an API call to the Members endpoint with the mc_eid as the value for the unique_email_id parameter.

For example, if your subscriber's link contains the tracking parameter mc_eid=18d3c1adfe, you can find that customer's email address by making a GET request to /lists/{list_id}/members?unique_email_id=18d3c1adfe

The mc_cid parameter is the Mailchimp ID for the campaign that generated the link.

The same record on the command line

Mailchimp says mc_eid can be turned back into the email address. I marked the two lines that say so in the sentences I pulled with curl from its developer page on 27 September 2026.

What I found Mailchimp's developer page describes mc_eid as "a unique Mailchimp ID that identifies the recipient's email address". According to the page, the sender can query the Members endpoint (the Mailchimp address that returns subscriber details) with this value. With that query the sender can find the person's email address. The example is mc_eid=18d3c1adfe. mc_cid is the ID of the campaign that generated the link.

Watch out *Your email address is not spelled out in mc_eid. The sender can still turn it back into the address. So opening or forwarding such a link can show that the click was yours.*

Keep in mind

Even in a separate work environment, check which account you are signed in to and which connection you are using.

01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27



Your turn

Your VPN is on. You are also signed in to your personal account. Is your identity hidden?

Answer: page 63



Puzzle

When I finished the Cover Your Tracks test, the results link in the address bar carried a parameter called `fpi_whorls`. It was 1,022 characters long once decoded (turned from its address-safe coding back into plain text). Had I posted that link on a forum, what would I have shared along with the address? Hint: run the test yourself and read the words after `fpi_whorls` in the address bar.

Answer: page 63



Check what you learned

You open an archived copy of the target's site. Can you assume the target will never see the visit?

Answer: page 63



Terms in this chapter

OPSEC	Operations security: protecting your identity, your work and your data during research.
Threat model	Thinking through, before you start, who could see what and how that could cause harm.
Passive collection	Here: reading data already held by another service.
Active collection	Here: sending a new request to the target.
Data minimization	Keeping only the data needed for the work.
Sock puppet / Browser fingerprint	A separate account used only for research / a trace, made of browser features such as language and screen size, that can set you apart.
Canary token / Tracking pixel	A hidden marker that may send a request to a server when it is opened or used. For example, a remote image or a marked file.

Sources

1. CSINT Research, WannaCry kill-switch domain records, captured 8 September 2026
<https://csintresearch.org/haber/wannacry-kill-switch>
2. Google, Incognito limits
<https://support.google.com/chrome/answer/9845881?hl=en>
3. Thinkst, Canarytokens
<https://docs.canarytokens.org/guide/>

Real cases

Each chapter in this part looks at one real event. I write the question, line up the public records with their dates and separate what each record shows from what it does not. Here you see the methods of the first five parts applied to real documents.

- 19 Did the queue clear when the Suez Canal reopened in 2021?
- 20 How long did Kīlauea's summit take to collapse in 2018?*(in the full edition)*
- 21 The red cells on the 2020 Beirut damage map*(in the full edition)*
- 22 Three dates on the WannaCry domain*(in the full edition)*
- 23 Two photo files in Vice's McAfee article*(in the full edition)*
- 24 The name on the BTK disk and the records that tested it*(in the full edition)*
- 25 What I could test in the Gilgo Beach bail filing*(in the full edition)*
- 26 Whose finding is each link of a 2026 phishing chain?*(in the full edition)*
- 27 What ties Bamako's 2025 shortages to the attacks on its roads? *(in the full edition)*

Did the queue clear when the Suez Canal reopened in 2021?

The canal reopened on 29 March 2021. By the Leth Agencies figure that NASA relays, 184 ships were still waiting the next day. In this chapter I test the reopening and the queue as two separate questions.

Event Gulf of Suez, Egypt, 23-29 March 2021

Question *When the canal reopened on 29 March 2021, had the queue of waiting ships also cleared?*

In this chapter

19.1 What happened?

19.2 The records I had

19.3 Small traces in the daytime image

19.4 The same gulf at night

19.5 What does the 29 March notice say?

19.1 What happened?

The container ship Ever Given ran aground in the Suez Canal on 23 March 2021. Passage between the Mediterranean and the Red Sea stopped, and ships waiting to enter the canal gathered at both ends. The European Space Agency (ESA) published a Sentinel-1 radar comparison on 26 March. It puts routine traffic on 21 March beside the blockage on 25 March.

On 29 March Ever Given was refloated and ships began moving through the canal again. I wanted to know whether the ships gathered at the entrance cleared that day too. Both questions belong to the same event, but their answers sit in different records. So I wrote them down separately from the start.




Figure 19.1 The canal reopened on 29 March, and 184 ships were still in line the next day. Keep those two days apart.

19.2 The records I had

For each record I wrote down two dates, the day it was published and the day its image was taken. NASA's article was published on 31 March, but the images inside it were taken between 1 February and 29 March.

NASA's file names carry a date too. The 202186 in `suezcanaoli_202186_lrg.jpg` means day 86 of 2021, which is 27 March. I turned the day numbers in all five file names into dates, and all five match the capture dates NASA gives.

The records, with publication and acquisition dates

Record	Published	Taken or reported
NASA Earth Observatory, daytime images (Landsat 8)	31 March 2021	11 and 27 March 2021
NASA Earth Observatory, night images (VIIRS)	31 March 2021	1 February, 27 and 29 March 2021
Suez Canal Authority notice	29 March 2021	Refloating of the ship
ESA, Sentinel-1 radar comparison	26 March 2021	21 and 25 March 2021




Figure 19.2 I applied the same steps to all five NASA file names. Find the day number from a file name on a calendar yourself and compare it with the capture date given.

19.3 **Small traces in the daytime image**

The two daytime images were taken on 11 and 27 March. First I wanted to be sure I was looking at the same place. The coastline and the narrowing northern end of the gulf match in both, so I applied the same crop to each.

On 27 March many more small, elongated traces stand out at sea. NASA identifies them as waiting ships. I did not count them. Clouds cover parts of the scene, and small traces lying side by side cannot always be told apart. Had I read a ship count off this image, I would have been pretending to count what I could not see.





Figure 19.3 The coastline and the northern end of the gulf match in both frames, so I put the box on the same pixels in both. In the 27 March frame there are many more small, long traces inside the box.

Image: NASA Earth Observatory / Lauren Dauphin, Joshua Stevens; USGS Landsat 8. NASA image use guidelines: may be used with credit. This does not mean NASA endorses the book.



11 March, inside the box



27 March, inside the box


Figure 19.4 I cut the inside of the box from NASA's full-resolution file. NASA reads these streaks as waiting ships. Because of the clouds and the traces that sit side by side, I did not count them.

Image: NASA Earth Observatory / Lauren Dauphin, Joshua Stevens; USGS Landsat 8. NASA image use guidelines: may be used with credit. This does not mean NASA endorses the book.

19.4 **The same gulf at night**

The night images come from VIIRS, an instrument on the Suomi NPP satellite that records light at night. I first found the city lights on land and then moved on to the points of light inside the gulf. Looking in that order makes the gathering of lights at sea easier to notice.

The daytime and night images have different resolutions, so I did not lay one over the other and match them pixel by pixel. Nor could I get a ship's name or cargo from a point of light. All I learned from the night image is that there were lights inside the gulf that night.



Night of 27 March 2021

Inside the box, 4 times

Figure 19.5 I found the city lights on land first, then looked inside the gulf. In the enlargement on the right I did not smooth the pixels, so each square is a pixel of the image. No ship name or count comes from the points of light.

Image: NASA Earth Observatory / Lauren Dauphin, Joshua Stevens; Suomi NPP VIIRS. NASA image use guidelines: may be used with credit. This does not mean NASA endorses the book.




Figure 19.6 All the night image gives is that there were lights in the gulf. A ship count or a ship's name needs another record.


19.5 **What does the 29 March notice say?**

The headline of the Suez Canal Authority's notice of 29 March announces that Ever Given was successfully refloated. Had I stopped at the headline, I might have thought the work was over. The body changes

that. According to it, the ship's heading had been corrected by 80% and manoeuvres would continue. Navigation would resume after full alignment.

In its article of 31 March, NASA writes that ships began moving through the canal again on 29 March and that clearing the waiting ships would take time. The same article, citing the canal service company Leth Agencies, reports that 184 ships were still waiting on 30 March. The notice and the article do not contradict each other.

They describe different moments of the same day. By the Leth Agencies figure, the queue had not cleared on 29 March. I took the number from NASA and did not verify it myself.



29 March 2021, Sentinel-2



Inside the box, 3 times

Figure 19.8 On the day of the notice, the gulf still holds many traces. Worldview gives only the day for this image, not the time. So I do not know whether it was taken before or after the ship was refloated.

Image: NASA Worldview (ESDIS); HLS S30, NASA LP DAAC; contains modified Copernicus Sentinel data 2021. NASA data with no use restrictions, credited to Worldview.

From the notice, summarised (29 March 2021)

Headline: Successful Refloating of EVER GIVEN
Heading corrected: 80%
Status: manoeuvres continue
Navigation: after full alignment

See the source list for the full notice.




Figure 19.7 The left column gives what each record shows; the dashed right column gives what may be read from it. No AIS data was used, so ship identities and an exact queue length cannot be derived.

What the records let me say

- The canal reopened on 29 March 2021. NASA's article gives that date. The canal authority's notice adds that passage would start once the ship was fully aligned.
- On 27 March there are many more traces in the gulf than on 11 March. NASA identifies them as waiting ships.
- The 29 March notice describes work still under way. By the Leth Agencies figure NASA relays, 184 ships were still waiting on 30 March.

What I cannot say

- How many ships were waiting, from the images. I did not count the traces. The figure of 184 above is Leth Agencies' count, as relayed by NASA.
- The name or cargo of any ship in the images.
- The day or hour when the queue fully cleared.
- The economic cost of the blockage, or who was responsible.

Methods used in this case

- Chapter 1, Write the question first (page 12)
- Chapter 2, Separate source and claim (page 24)
- Chapter 6, Look at a photo three ways (in the full edition)
- Chapter 7, Where? When? (in the full edition)

Your turn

Suppose a news site had written on the evening of 29 March that "traffic at Suez is back to normal". Which record would you use to test that sentence? Which of the records in this chapter would fall short?

Answer: page 64

Sources

1. NASA Earth Observatory, Traffic Jam on the Suez Canal (31 March 2021). Daytime and night images, their dates and NASA's reading.
<https://science.nasa.gov/earth/earth-observatory/traffic-jam-on-the-suez-canal-148114/>
2. Suez Canal Authority, Successful Refloating of EVER GIVEN (29 March 2021). Headline and body read together.
https://www.suezcanal.gov.eg/English/MediaCenter/News/Pages/nav_29-03-2021.aspx
3. ESA, Suez Canal traffic jam seen from space (26 March 2021). A separate radar observation. Its images are not used in this book.
https://www.esa.int/ESA_Multimedia/Images/2021/03/Suez_Canal_traffic_jam_seen_from_space
4. NASA, Images and Media Usage Guidelines. Terms for using NASA images.
<https://www.nasa.gov/nasa-brand-center/images-and-media/>




Figure 19.9 *Small traces in the daytime image.* The same crop on two dates: 11 March on the left, 27 March 2021 on the right. The box marks the same place in both. The Python session turns the day number in NASA's file names into a date. Images: NASA Earth Observatory, USGS Landsat 8.




Figure 19.10 *Small traces in the daytime image.* A working copy of the 27 March image. The rings mark a few of the traces NASA reads as waiting ships; nothing shows under the cloud in the blue box. I did not count the traces. Image: NASA Earth Observatory, USGS Landsat 8.




Figure 19.11 *The same gulf at night.* Three nights with the same crop: 1 February, 27 March and 29 March 2021. The boxes mark the northern end of the gulf, where lights gather on 27 and 29 March. Images: NASA Earth Observatory, Suomi NPP VIIRS.

Answer key

1. Write the question first

How would you narrow the task “Collect information about the Suez Canal”?

For example, you could ask “Were ships waiting to enter the canal in the Gulf of Suez on 27 March 2021?” Both the place and the date are fixed. You can look for the answer in a record, such as NASA’s satellite image of that day. Any other question that is just as narrow would also work.

symbolics.com is known as the first .com name ever registered. Its RDAP record (its registration data, at rdap.verisign.com/com/v1/domain/symbolics.com) gives a registration time of 1985-03-15T05:00:00Z. The expiry time is 2027-03-16T04:00:00Z. The T separates the date from the time, and the Z at the end means the time is in UTC. Why does the same record say 05:00 in one place and 04:00 in the other? Hint: think about daylight saving time.

Both are midnight in New York. According to the IANA time zone database, US daylight saving time began on the last Sunday of April from 1976 to 1986. So on 15 March 1985 it had not yet begun. New York was on UTC-5 and midnight was 05:00 UTC. Since 2007 it has begun on the second Sunday of March. So on 16 March 2027 New York will be on daylight saving time, UTC-4.

Midnight there will be 04:00 UTC. [google.com](https://www.google.com), registered in September, also shows 04:00 UTC, because daylight saving time was in force on that day too.

The task asks whether a service outage affects a company. Your first list has 500 staff names. What do you do?

First I check which service the company uses. A list of 500 staff names does not answer the question, and I do not keep personal data the question does not need.

2. Separate source and claim

Does F6 in the Admiralty system mean “false”?

No. Whoever writes F6 records that neither the source's reliability nor the information's credibility could be judged. You cannot call something false just because it is unknown.

In a column dated 4 May 2009, the Guardian's readers' editor wrote about the invented quote. According to the column, the second time it was added to Wikipedia it stayed for only six minutes. The third time it stayed for about 25 hours. According to the Maurice Jarre page history for 30 and 31 March 2009, which insertion stayed for six minutes? Hint: on the history page from task 2, look at the plus or minus byte count on each line. Then use "prev" to check the lines that add or remove the quote.

The third. The history shows three insertions. The first was added on 30 March at 02:29 and removed at 11:51. The second was added on 30 March at 14:13. It was removed at 15:07 on 31 March with the note "no unsourced quotes", so it stayed about 25 hours. The third was added on 31 March at 17:03 and reverted at 17:09, so it stayed six minutes.

The column's order does not match the page history. The column also says the quote was deleted twice on 30 March, but the history shows only one removal that day.

Three articles use the same anonymous post. How do you count the sources?

There are three publications but only one origin for the information. I do not count the repetition as new evidence. I look for independent confirmation.

3. Prepare before searching

Your VPN is on. You are also signed in to your personal account. Is your identity hidden?

No. A VPN changes your exit IP address, but it does not hide the account you signed in to. The site sees that account.

When I finished the Cover Your Tracks test, the results link in the address bar carried a parameter called `fpi_whorls`. It was 1,022 characters long once decoded (turned from its address-safe coding back into plain text). Had I posted that link on a forum, what would I have shared along with the address? Hint: run the test yourself and read the words after `fpi_whorls` in the address bar.

Part of my browser's fingerprint. The parameter is JSON (a text format that stores data as names and values) embedded in the address. Its fields are plugins (the plugin list), hardware_concurrency (the number of CPU cores) and audio (the audio fingerprint). It also has canvas_hash_v2 and webgl_hash_v2 (hashes of the drawing tests). My name and IP address are not in it. But the page says the canvas result tends to stay consistent enough on one machine to identify a user.

You open an archived copy of the target's site. Can you assume the target will never see the visit?

No. An archived page may load some images from the live site. I do not assume there is no contact with the target until I have checked which requests go where.

19. Did the queue clear when the Suez Canal reopened in 2021?

Suppose a news site had written on the evening of 29 March that "traffic at Suez is back to normal". Which record would you use to test that sentence? Which of the records in this chapter would fall short?

The sentence is about traffic in the canal. To test it, you need a dated record that counts the ships passing through the canal on the days after 29 March. Later notices from the canal authority would do, for example, or dated ship-tracking (AIS) data. NASA relays a Leth Agencies figure of 184 ships still waiting on 30 March.

That figure already contradicts the sentence. But it is secondhand, and it needs a primary record to confirm it. The 29 March notice is no use, because it records that manoeuvres were still going on. The 29 March night image is no use either. It gives no time of capture, and nothing in it ties the lights to particular ships.

Images and rights

Diagrams and tables were drawn for this book from the text they explain. Photographs and satellite images are public records, used with the credit their owners ask for. The screenshots in the field tasks are cropped views of public pages, quoted for analysis with their source and date; the terminal images show commands I ran and their real output. None of the images was generated by artificial intelligence.

1. Cover: the Gulf of Suez, 27 March 2021, Landsat 8. NASA identifies the small traces at sea as waiting ships. The image is one of the records examined in Chapter 19. NASA Earth Observatory / Lauren Dauphin, Joshua Stevens; USGS Landsat 8.
2. Chapter 19, work screen: The same crop on two dates: 11 March on the left, 27 March 2021 on the right. The box marks the same place in both. The Python session turns the day number in NASA's file names into a date. Images: NASA Earth Observatory, USGS Landsat 8.
3. Chapter 19, work screen: A working copy of the 27 March image. The rings mark a few of the traces NASA reads as waiting ships; nothing shows under the cloud in the blue box. I did not count the traces. Image: NASA Earth Observatory, USGS Landsat 8.
4. Chapter 19, work screen: Three nights with the same crop: 1 February, 27 March and 29 March 2021. The boxes mark the northern end of the gulf, where lights gather on 27 and 29 March. Images: NASA Earth Observatory, Suomi NPP VIIRS.
5. Chapter 1, in the field: The 1990-09-17 registration date, boxed in red at the bottom of IANA's .tr page. IANA, screenshot 27.09.2026.
6. Chapter 1, in the field: The TÜİK bulletin on my phone. The line boxed in red is the release date, 9 February 2026, 10:00. Türkiye İstatistik Kurumu (TÜİK), screenshot 27.09.2026.
7. Chapter 3, in the field: "unique", the 348,199 browsers and the 18.41 bits, boxed in red in the Cover Your Tracks result. EFF Cover Your Tracks (CC BY), screenshot 27.09.2026.

The full edition

This file holds chapters 1, 2, 3 and 19 of the book's 27. The full edition has 18 method chapters and 9 real cases retraced from public records. At the back are an appendix, "Public records in Türkiye and the law", the answers to every exercise, a glossary and the full source list. It comes as a PDF.

Part 1: First, how to think

- 01 Write the question first *(in this file)*
- 02 Separate source and claim *(in this file)*
- 03 Prepare before searching *(in this file)*

Part 2: Find, save, check

- 04 Narrow the search
- 05 Keep what you find
- 06 Look at a photo three ways
- 07 Where? When?
- 08 Put the rows in order

Part 3: People and companies

- 09 Separate account and person
- 10 Find the record behind the brand

Part 4: Internet records

- 11 Name, address and DNS
- 12 Read the name in the certificate
- 13 Which name do email checks verify?
- 14 Whose server? What evidence?

Part 5: Finish the work

- 15 Privacy and keys
- 16 Learn a tool by its job
- 17 Work with a small file
- 18 Write what you know

Part 6: Real cases

- 19 **Did the queue clear when the Suez Canal reopened in 2021?** *(in this file)*
- 20 How long did Kīlauea’s summit take to collapse in 2018?
- 21 The red cells on the 2020 Beirut damage map
- 22 Three dates on the WannaCry domain
- 23 Two photo files in Vice’s McAfee article
- 24 The name on the BTK disk and the records that tested it
- 25 What I could test in the Gilgo Beach bail filing
- 26 Whose finding is each link of a 2026 phishing chain?
- 27 What ties Bamako’s 2025 shortages to the attacks on its roads?

For the full edition: en.csintresearch.org/gosi-book

About the author

Ahmet Göker is an open-source researcher and threat analyst. He publishes research files, tools and study notes at CSINT Research. He also posts step-by-step videos of real cases on YouTube at @shadowintel_. Ahmet likes puzzles and true-crime documentaries. He is most interested in how one small detail can crack a case that went unsolved for years.

`csintresearch.org`
`youtube.com/@shadowintel_`