

FOR RELEASE MAY 17, 2024

When Online Content Disappears

38% of webpages that existed in 2013 are no longer accessible a decade later

BY *Athena Chapekis, Samuel Bestvater, Emma Remy and Gonzalo Rivero*

FOR MEDIA OR OTHER INQUIRIES:

Aaron Smith, Director, Data Labs
Sogand Afkari, Communications Manager

202.419.4372
www.pewresearch.org

RECOMMENDED CITATION

Pew Research Center, May 2024, "When Online Content Disappears"

How we did this

Pew Research Center conducted the analysis to examine how often online content that once existed becomes inaccessible. One part of the study looks at a representative sample of webpages that existed over the past decade to see how many are still accessible today. For this analysis, we collected a sample of pages from the [Common Crawl](#) web repository for each year from 2013 to 2023. We then tried to access those pages to see how many still exist.

A second part of the study looks at the links on existing webpages to see how many of those links are still functional. We did this by collecting a large sample of pages from government websites, news websites and the online encyclopedia [Wikipedia](#).

We identified relevant news domains using data from the audience metrics company [comScore](#) and relevant government domains (at multiple levels of government) using data from [get.gov](#), the official administrator for the .gov domain. We collected the news and government pages via Common Crawl and the Wikipedia pages from an archive maintained by the [Wikimedia Foundation](#). For each collection, we identified the links on those pages and followed them to their destination to see what share of those links point to sites that are no longer accessible.

A third part of the study looks at how often individual posts on social media sites are deleted or otherwise removed from public view. We did this by collecting a large sample of public tweets on the social media platform X (then known as Twitter) in real time using the Twitter Streaming API. We then tracked the status of those tweets for a period of three months using the Twitter Search API to monitor how many were still publicly available.

Refer to the report [methodology](#) for more details.

When Online Content Disappears

38% of webpages that existed in 2013 are no longer accessible a decade later


The internet is an unimaginably vast repository of modern life, with hundreds of billions of indexed webpages. But even as users across the world rely on the web to access books, images, news articles and other resources, this content sometimes disappears from view.

A new Pew Research Center analysis shows just how fleeting online content actually is:

- A quarter of all webpages that existed at one point between 2013 and 2023 are no longer accessible**, as of October 2023. In most cases, this is because an individual page was deleted or removed on an otherwise functional website.
- For older content, this trend is even starker.** Some 38% of webpages that existed in 2013 are not available today, compared with 8% of pages that existed in 2023.

38% of webpages from 2013 are no longer accessible

% of links from each year that are no longer accessible as of October 2023



Source: Pew Research Center analysis of a random selection of URLs collected by the Common Crawl web repository (n=999,989) and checked using page and DNS response codes. Web pages defined as inaccessible if they returned a status code of 204, 400, 404, 410, 500, 501, 502, 503, 523 or did not return a valid status code.

“When Online Content Disappears”

PEW RESEARCH CENTER

This “digital decay” occurs in many different online spaces. We examined the links that appear on government and news websites, as well as in the “References” section of Wikipedia pages as of spring 2023. This analysis found that:

- 23% of news webpages contain at least one broken link, as do 21% of webpages from government sites.** News sites with a high level of site traffic and those with less are about equally likely to contain broken links. Local-level government webpages (those belonging to city governments) are especially likely to have broken links.

- **54% of Wikipedia pages contain at least one link in their “References” section that points to a page that no longer exists.**

To see how digital decay plays out on social media, we also collected a real-time sample of tweets during spring 2023 on the social media platform X (then known as Twitter) and followed them for three months. We found that:

- **Nearly one-in-five tweets are no longer publicly visible on the site just months after being posted.** In 60% of these cases, the account that originally posted the tweet was made private, suspended or deleted entirely. In the other 40%, the account holder deleted the individual tweet, but the account itself still existed.
- **Certain types of tweets tend to go away more often than others.** More than 40% of tweets written in Turkish or Arabic are no longer visible on the site within three months of being posted. And tweets from accounts with the default profile settings are especially likely to disappear from public view.

How this report defines inaccessible links and webpages

There are many ways of defining whether something on the internet that used to exist is now inaccessible to people trying to reach it today. For instance, “inaccessible” could mean that:

- The page no longer exists on its host server, or the host server itself no longer exists. Someone visiting this type of page would typically receive a variation on the [“404 Not Found”](#) server error instead of the content they were looking for.
- The page address exists but its content has been changed – sometimes dramatically – from what it was originally.
- The page exists but certain users – such as those with blindness or other visual impairments – might find it difficult or impossible to read.

For this report, we focused on the first of these: pages that no longer exist. The other definitions of accessibility are beyond the scope of this research.

Our approach is a straightforward way of measuring whether something online is accessible or not. But even so, there is some ambiguity.

First, there are dozens of status codes indicating a problem that a user might encounter when they try to access a page. Not all of them definitively indicate whether the page is permanently defunct or just temporarily unavailable. Second, for security reasons, many sites actively try to prevent the sort of automated data collection that we used to test our full list of links.

For these reasons, we used the most conservative estimate possible for deciding whether a site was actually accessible or not. We counted pages as inaccessible only if they returned one of nine error codes that definitively indicate that the page and/or its host server no longer exist or have become nonfunctional – regardless of how they are being accessed, and by whom. The full list of error codes that we included in our definition are in the [methodology](#).

Here are some of the findings from our analysis of digital decay in various online spaces.

Webpages from the last decade

To conduct this part of our analysis, we collected a random sample of just under 1 million webpages from the archives of [Common Crawl](#), an internet archive service that periodically collects snapshots of the internet as it exists at different points in time. We sampled pages collected by Common Crawl each year from 2013 through 2023 (approximately 90,000 pages per year) and checked to see if those pages still exist today.

We found that 25% of all the pages we collected from 2013 through 2023 were no longer accessible as of October 2023. This figure is the sum of two different types of broken pages: 16% of pages are individually inaccessible but come from an otherwise functional root-level domain; the other 9% are inaccessible because their entire root domain is no longer functional.

Not surprisingly, the older snapshots in our collection had the largest share of inaccessible links. Of the pages collected from the 2013 snapshot, 38% were no longer accessible in 2023. But even for pages collected in the 2021 snapshot, about one-in-five were no longer accessible just two years later.


Links on government websites

We sampled around 500,000 pages from government websites using the Common Crawl March/April 2023 snapshot of the internet, including a mix of different levels of government (federal, state, local and others). We found every link on each page and followed a random selection of those links to their destination to see if the pages they refer to still exist.

Across the government websites we sampled, there

Around 1 in 5 government webpages contain at least one broken link

Broken links on government websites, by level of government



Note: Links are defined as broken if the root domain is not available in a DNS server, or if the site returned an error indicating the page is not available.

Source: Pew Research Center analysis of links on 500,000 government webpages collected from Common Crawl March/April 2023 crawl. 10% of the links on each page were sampled, for a total of n=4,200,000 links analyzed.

“When Online Content Disappears”

PEW RESEARCH CENTER

were 42 million links. The vast majority of those links (86%) were internal, meaning they link to a different page on the same website. An explainer resource on the IRS website that links to other documents or forms on the IRS site would be an example of an internal link.

Around three-quarters of government webpages we sampled contained at least one on-page link. The typical (median) page contains 50 links, but many pages contain far more. A page in the 90th percentile contains 190 links, and a page in the 99th percentile (that is, the top 1% of pages by number of links) has 740 links.

Other facts about government webpage links:

- The vast majority go to secure HTTP pages (and have a URL starting with “https://”)
- 6% go to a static file, like a PDF document
- 16% now redirect to a different URL than the one they originally pointed to

When we followed these links, we found that 6% point to pages that are no longer accessible. Similar shares of internal and external links are no longer functional.

Overall, 21% of all the government webpages we examined contained at least one broken link. Across every level of government we looked at, there were broken links on at least 14% of pages; city government pages had the highest rates of broken links.

Links on news websites

For this analysis, we sampled 500,000 pages from 2,063 websites classified as “News/Information” by the audience metrics firm comScore. The pages were collected from the Common Crawl March/April 2023 snapshot of the internet.

Across the news sites sampled, this collection contained more than 14 million links pointing to an outside website.¹ Some 94% of these pages contain at least one external-facing link. The median page contains 20 links, and pages in the top 10% by link count have 56 links.

Like government websites, the vast majority of these links go to secure HTTP pages (those with a URL beginning with “https://”). Around 12% of links on these news sites point to a static file, like a PDF document. And 32% of links on news sites redirected to a different URL than the one they originally pointed to – slightly less than the 39% of external links on government sites that redirect.


When we tracked these links to their destination, we found that 5% of all links on news site pages are no longer accessible. And 23% of all the pages we sampled contained at least one broken link.

Broken links are about as prevalent on the most-trafficked news websites as they are on the least-trafficked sites. Some 25% of pages on news websites in the top 20% by site traffic have at least one broken link. That is nearly identical to the 26% of sites in the bottom 20% by site traffic.

23% of news webpages have at least one broken link

Broken links on news websites, by site traffic ranking

■ % of **pages with at least one** broken link
 ■ % of **all links** that are broken



Note: Links defined as broken if the root domain is not available in a DNS server, or if the site returned an error indicating the page is not available.

Source: Pew Research Center analysis of links on 500,000 news webpages (n=2,063) collected from Common Crawl March/April 2023 crawl. 50% of all links on these pages were sampled, for a total of n=7,089,514 links analyzed. Sites and traffic quintiles identified using data from comScore.

“When Online Content Disappears”

PEW RESEARCH CENTER

¹ For our analysis of news sites, we did not collect or check the functionality of internal-facing on-page links – those that point to another page on the same root domain.

Reference links on Wikipedia

For this analysis, we collected a random sample of 50,000 English-language Wikipedia pages and examined the links in their “References” section. The vast majority of these pages (82%) contain at least one reference link – that is, one that directs the reader to a webpage other than Wikipedia itself.

In total, there are just over 1 million reference links across all the pages we collected. The typical page has four reference links.

The analysis indicates that 11% of all references linked on Wikipedia are no longer accessible. On about 2% of source pages containing reference links, every link on the page was broken or otherwise inaccessible, while another 53% of pages contained at least one broken link.


Posts on Twitter

For this analysis, we collected nearly 5 million tweets posted from March 8 to April 27, 2023, on the social media platform X, which at the time was known as Twitter. We did this using Twitter’s Streaming API, collecting 3,000 public tweets every 30 minutes in real time. This provided us with a representative sample of all tweets posted on the platform during that period. We monitored those tweets until June 15, 2023, and checked each day to see if they were still available on the site or not.

At the end of the observation period, we found that **18% of the tweets from our initial collection window were no longer publicly visible on the site.** In a majority of cases, this was because the account that originally posted the tweet was made private, suspended or deleted entirely. For the remaining tweets, the account that posted the tweet was still visible on the site, but the individual tweet had been deleted.

Around 1 in 5 tweets disappear from public view within months

% of tweets posted March 8-April 27, 2023, that were available/unavailable as of June 15, 2023



Source: Pew Research Center analysis of 4.8 million tweets posted March 8-April 23, 2023. Analysis of tweet- and account-level characteristics based on a random sample of those tweets (n=148,494). Tweets collected using Twitter Streaming API and monitored using Twitter Search API. “When Online Content Disappears”

PEW RESEARCH CENTER

Which tweets tend to disappear?

Tweets were especially likely to be deleted or removed over the course of our collection period if they were:


- **Written in certain languages.** Nearly half of all the Turkish-language tweets we collected – and a slightly smaller share of those written in Arabic – were no longer available at the end of the tracking period.
- **Posted by accounts using the site’s default profile settings.** More than half of tweets from accounts using the default profile image were no longer available at the end of the tracking period, as were more than a third from accounts with a default bio field. Tweets from these accounts tend to disappear because the entire account has been deleted or made private, as opposed to the individual tweet being deleted.
- **Posted by unverified accounts.**

We also found that removed or deleted tweets tended to come from newer accounts with relatively few followers and modest activity on the site. On average, tweets that were no longer visible on the site were posted by accounts around eight months younger than those whose tweets stayed on the site.

And when we analyzed the types of tweets that were no longer available, we found that retweets, quote tweets and original tweets did not differ much from the overall average. But replies were relatively unlikely to be removed – just 12% of replies were inaccessible at the end of our monitoring period.

Inaccessible tweets often come from accounts with default profile settings

% of tweets posted March 8-April 27, 2023, that were no longer publicly visible on the site as of June 15, 2023



Source: Pew Research Center analysis of 4.8 million tweets posted March 8-April 23, 2023. Analysis of tweet- and account-level characteristics based on a random sample of those tweets (n=148,494). Tweets collected using Twitter Streaming API and monitored using Twitter Search API. “When Online Content Disappears”

PEW RESEARCH CENTER

Most tweets that are removed from the site tend to disappear soon after being posted. In addition to looking at how many tweets from our collection were still available at the end of our tracking period, we conducted a [survival analysis](#) to see how long these tweets tended to remain available. We found that:

- 1% of tweets are removed within one hour
- 3% within a day
- 10% within a week
- 15% within a month

Put another way: Half of tweets that are eventually removed from the platform are unavailable within the first six days of being posted. And 90% of these tweets are unavailable within 46 days.

Tweets don't always disappear forever, though. Some 6% of the tweets we collected disappeared and then became available again at a later point. This could be due to an account going private and then returning to public status, or to the account being suspended and later reinstated. Of those “reappeared” tweets, the vast majority (90%) were still accessible on Twitter at the end of the monitoring period.

Acknowledgments

This report is a collaborative effort based on the input and analysis of the following individuals:

Primary Researchers

Athena Chapekis, *Data Science Analyst*
Samuel Bestvater, *Computational Social Scientist*
Emma Remy, *Former Data Science Analyst*
Gonzalo Rivero, *Former Associate Director, Data Labs*

Research Team

Aaron Smith, *Director, Data Labs*
Brian Broderick, *Senior Data Engineer*
Galen Stocking, *Senior Computational Social Scientist*
Regina Widjaya, *Computational Social Scientist*
Meltem Odabaş, *Former Computational Social Scientist*

Editorial and Graphic Design

Alissa Scheller, *Senior Information Graphics Designer*
Anna Jackson, *Editorial Assistant*

Communications and Web Publishing

Sogand Afkari, *Communications Manager*
Janakee Chavda, *Assistant Digital Producer*

In addition, the project benefited greatly from feedback by Jeff Diamant, Jenn Hatfield, Monica Anderson and Lee Rainie of Pew Research Center.

Methodology

Collection and analysis of Twitter data

Twitter analysis in this report is based on 4.8 million tweets collected from March 8 to April 27, 2023. This process involved collecting batches of 3,000 new tweets every 30 minutes over the duration of the collection period using the Twitter Streaming API. This resulted in a sample of tweets created at different times and days over a number of weeks.

We regularly monitored the status of those tweets starting March 15 and ending June 15, 2023. Each day during the monitoring period, we looked up all collected tweets using the Twitter Search API. We collected the most recent engagement metrics for those tweets, as well as a status code indicating whether each tweet was still publicly available on the site or not.

Tweets were classified as unavailable if they returned a status code of “Not Found” (indicating the tweet itself had been deleted) or “Authorization Error” (indicating it was inaccessible because the account itself had been deleted or made private by the user or suspended by Twitter itself). Because we monitored the status of all collected tweets over the duration of the monitoring period, we were able to identify tweets that became visible again after previously being unavailable.

In addition to examining attrition rates using the full sample of 4.8 million tweets, we selected tweets from a random sample of 100,000 users, resulting in a sample of 148,494 tweets from our original collection, and gathered detailed information about those tweets and the users who posted them. These included details such as the language the tweet was written in; whether the bio field or profile picture of the account had been updated from the site defaults; the age of the account; and whether the account is verified. This subsample is used in the analysis of what types of tweets tend to be removed from the site.

Data collection for World Wide Web websites, government websites and news websites

To examine attrition on the broader internet, we collected three samples of web crawl data from [Common Crawl](#), a nonprofit organization that collects and maintains a large web archive. This data has been collected monthly starting in 2008. We used this archive to create a historical sample of URLs from the broader internet dating back to 2013, as well as contemporaneous snapshots of pages from government entities and news websites.

Sample of URLs from the broader internet

This analysis is based on random samples of URLs from crawls conducted from 2013 to 2023, using year as a stratifying variable. We used the March/April crawl where possible and the closest available date range for years in which a March/April crawl was not conducted. This resulted in a full sample of 1 million pages – approximately 91,000 pages each year from 2013 to 2023 – that were known to have existed at the time they were collected by Common Crawl.

We then looked at whether these pages were still available in fall 2023 using the procedure described below. These checks were performed in several stages, running Oct. 12-Nov. 6, 2023.

Sample of government website URLs

This analysis is based on a random sample of 500,000 pages with a .gov domain, stratified by domain and level of government. We collected these pages from the Common Crawl MAIN-2023-14 crawl conducted March/April 2023.

Each page was assigned to a level of government (Federal – Executive; State; City; County; Federal – Legislative; Federal – Judicial; Tribal; Independent Intrastate; and Interstate) using <https://get.gov>, the official administrator for the .gov top-level domain. We retrieved [the dataset used for this analysis](#) Aug. 22, 2023.

This resulted in a sample with the following breakdown of domains and levels of government:

Sample of government website URLs

Level of government	Total domains	Total pages	Average pages per domain
Federal – Executive	659	137,717	209
State	579	93,956	162
City	2403	71,162	30
County	818	57,267	70
Federal – Legislative	49	45,686	932
Federal – Judicial	6	38,683	6,447
Tribal	108	28,926	268
Independent Intrastate	92	21,797	237
Interstate	16	4,806	300

Source: Pew Research Center analysis of links on 500,000 government webpages collected from Common Crawl March/April 2023 crawl. “When Online Content Disappears”

PEW RESEARCH CENTER

For each of the 500,000 pages collected, we selected a random sample of 10% of all links (internal as well as external) found on that page. This resulted in a total of 4,179,313 links. We then looked at whether the pages these links point to were still available.

Sample of news website URLs

The analysis of news websites is based on a list of 2,063 domains categorized as “News/Information” by the measurement and audience metrics company comScore. We divided these domains into quintiles based on comScore site traffic for Q4 2022 and sampled 500,000 total pages from these domains using site traffic quintiles as a stratifying variable.

This resulted in a sample with the following breakdown of domains:

Sample of news website URLs

Q4 2022 web traffic quintile	Total domains	Total pages	Average pages per domain
Quintile 1 (lowest traffic)	398	100,000	251
Quintile 2	405	100,000	247
Quintile 3	421	100,000	238
Quintile 4	423	100,000	236
Quintile 5 (highest traffic)	416	100,000	240

Source: Pew Research Center analysis of links on 500,000 news webpages (n=2,063) collected from Common Crawl March/April 2023 crawl. Sites and traffic quintiles identified using data from comScore. “When Online Content Disappears”

PEW RESEARCH CENTER

We selected a 50% simple random sample of all the 7,089,514 links that appeared on these pages, excluding any internal links (those that point within the same host domain). We then looked at whether the pages these links point to were still available.

Data collection for Wikipedia source links

We sampled 50,000 pages from the list of all titles in the [English Wikipedia May 20, 2023, snapshot](#) on Sept. 20, 2023. As some pages have multiple titles in the list of all titles, but refer to the same page (for instance, “[UK](#)” and “[United Kingdom](#)”), we followed redirects to eliminate duplicate titles for the same page. Between the snapshot and our collection, 50 pages were removed; our analysis is based on the remaining 49,950 pages.

Our analysis evaluated all external links (that is, links pointing to non-Wikipedia domains) from the “References” section of all the pages in the sample as of Oct. 10-11, 2023, using the same definition of link and procedure described above.

Evaluating the status of pages and links

We categorized links as alive or dead using the response code from the page. A page was classified as inaccessible if the domain was not available in a DNS server or if the server returned one of the following error codes indicating the content was not available:

- 204 No Content
- 400 Bad Request
- 404 Not Found
- 410 Gone
- 500 Internal Server Error
- 501 Not Implemented
- 502 Bad Gateway
- 503 Service Unavailable
- 523 Origin Is Unreachable

Pages were considered accessible in all other cases – including ambiguous situations in which we could not guarantee that the content exists, like soft 404 pages or timeouts not caused by the DNS.

We evaluated links in four rounds. In the first round (Oct. 12 to Oct. 15), we evaluated whether links were functional by following them using the requests library in Python, allowing for pages to timeout after one second. In this round, we recorded the initial status code and final status code after redirects, if applicable.

For the pages that did not return a 200 OK status code, we did a second round of evaluations (Oct. 16 to Oct. 17) in which we collected the status code using randomized browser headers from the library `fake_headers`.

A third round (Oct. 27 to Oct. 28) rechecked pages that did not successfully resolve to any status code and for pages that returned a 429 (“Too many requests”) status code, with an additional timeout of three seconds.

In the final round (Nov. 6), we looked up the pages that did not return any result in a DNS server using the `dnspython` module in Python allowing for a three-second timeout.

Definition of links

We identified hyperlinks from the HTML code of the websites by looking at all <a> tags that included a href attribute. We limited our attention to hyperlinks that used the HTTP or HTTPS protocol. Pages frequently use relative links that do not include the specification of the scheme and domain of the site in the definition. In those cases, we restricted our attention to those that referred to subdomains or paths (i.e., that started with a backslash /) and discarded hyperlinks defined by anchors (i.e., that started with a pound sign #).

Whenever a page used a relative link, we tried reconstructing the absolute URL by prepending the domain information. In our analyses, these reconstructed URL were treated as any other URL during our analyses.