



Almost everything UCMR 5 measured is below the line

Labs report a PFAS result only when it reaches a minimum reporting level. Anything lower is written down as "< MRL": less than the line, not zero. Of the 1,925,302 PFAS results from 10,313 water systems, this is how they fall.



Each dot is about 1,925 results. All 29 PFAS, January 2023 to June 2026, water entering distribution systems.

The six PFAS with 2024 limits




For PFOA and PFOS the reporting level, 4 ppt, is the same as the federal limit. So the data can show a result at the limit, but cannot tell 0 ppt from 3.9 ppt, or show whether a system is under the rule's 2 ppt trigger level. The goal the rule sets for both is zero.

The rule counts non-detects as zero


The 2024 rule uses zero for any result below its quantitation level, which for these PFAS equals the UCMR 5 reporting level, "solely to calculate the running annual average" (40 CFR 141.903(f)(1)(iv)). EPA's Data Summary does not say how it treated non-detects, but its count reproduces exactly with zeros. Zero is the lowest value a non-detect can have, so every average built on it is a minimum.

Two results, three averages



An illustrative ground-water entry point, sampled twice for PFOA. The data place its true average between 3.1 and 5.1 ppt. Whether it counts as over the limit depends entirely on the value given to the non-detect.


The same choice, across the country



Only averages change; every other rule is fixed. EPA, with its own split of system sizes, reports 7.8%; the same weighting here gives 7.6% with zeros. Half the reporting level is shown because it is a common shortcut, not because it is more likely.

Where the data can decide, and where they can't

10,129 water systems have complete sample sets. Compared with the 2024 limits by EPA's method, each falls into one of three groups, depending on whether the answer changes when the non-detects move from zero up to the line.



1,136

over a limit on the measured values alone
54.5 million people

492

could be over or under:
it depends on where the non-detects really are
22.6 million people


8,501

under, even if every non-detect sat just below the line

Each square is 10 systems. For PFOA and PFOS alone: 1,127 over, 476 undecided, 8,526 under. People are SDWIS "population served".

The Hazard Index can't be confirmed under 1

The index adds four PFAS, each divided by a benchmark. At 23,023 of 26,703 entry points (86%) none of the four was detected. Yet with each just under its reporting level, the index could still reach 1.20.



PFBS adds 0.0015 (under $3 \div 2,000$). EPA's comparison counts the index only where two or more of the four were detected and the average is 1.5 or more: 80 systems with zeros, up to 235 with every non-detect at the line.

The top of the distribution is measured. The middle isn't.


Percentiles above the share of non-detects are real measurements. Everything else about the typical level is a range, because no one knows where under the line the non-detects sit.



All results nationally, log scale. Red: Kaplan–Meier percentiles, with 95% intervals from 2,000 bootstrap resamples of whole systems. Hatched: the range the national mean must lie in, from non-detects at zero to non-detects at the line (the Kaplan–Meier mean).

Where a fitted distribution can fill in the gap


Regression on order statistics (ROS) fits a lognormal curve to the detected values and estimates what lies under the line. It needs at least 8 detects and at least 20% detected. No PFAS meets that nationally; 28 of 1,624 state and PFAS combinations do, including these 7 for PFOA and PFOS.



Hatched: substitution bounds. Diamond and line: ROS mean with its 95% interval. In every case the estimate falls inside the bounds and is far narrower. Substituting half the reporting level gives a number with no statistical basis (Helsel 2006).


Big systems were all tested. Small ones were sampled.

UCMR 5 required every system serving 3,300 or more people to test, plus a nationally representative sample of 800 smaller ones (86 FR 73131, 2021). Dark squares are tested systems, light squares the rest, drawn to scale.



For larger systems the counts are close to a census: the 662 large and 419 medium systems EPA found over a limit are counts, not estimates. For small systems only a national rate exists: 55 of 775, or 7.1% (5.5–9.1%). They are 85% of all systems, so that rate drives EPA's 7.8% for the whole country. Not tested at all: 56,144 systems serving 34.0 million people.

Too few small systems per state for a state rate




28 of the 56 states and territories had 12 or fewer small systems tested; 22 had under 10 and 3 had none. At the national 7.1%, a rate from 12 systems would carry a 95% margin of about ± 16 points; even Texas, with 112, would be ± 5 . Sizes here use 2026 populations; EPA drew its sample using 2021 figures, which is why 987 systems count as small here.

Population is what each system reports as "population served". A system counts in full if any one entry point is flagged, and schools and workplaces count students and workers. None of it measures what anyone drinks.


Source water changes what labs can detect

Ground water, surface water, ground water under the direct influence of surface water (GU) and mixed sources are kept apart. Ground-water entry points were sampled twice and the others four times, so the per-sample measure is the fair comparison.



Lines: 95% intervals (2,000 bootstrap resamples of whole systems). Surface-water samples detect some PFAS more often than ground water (25.9% vs 16.3%), but their entry points are not more often over a limit (6.9% vs 8.0%); more sit just under it, undecided (6.4% vs 2.4%). Exploratory; Benjamini–Hochberg q in the output tables.

Who is in the small-system sample



EPA allocated its small-system sample in proportion to population served, so systems serving 500 or fewer people are 72% of small systems but 23% of those tested. Among tested systems their rate was 8.7% (218 systems) against 6.4% for 501–3,299 (738); a difference within sampling error, but it changes the projection (next page).


Rate over a limit in each stratum

People served	Ground water		Surface water		GU	
More than 10,000	17.1%	1,660 of 1,702 tested	14.2%	2,767 of 2,904 tested	22.6%	53 of 56 tested
3,300 to 10,000	8.0%	2,602 of 2,881 tested	8.0%	2,024 of 2,251 tested	16.4%	67 of 71 tested
501 to 3,299	7.1%	522 of 11,767 tested	5.0%	202 of 3,633 tested	0.0%	14 of 215 tested, few
500 or fewer	8.0%	187 of 37,383 tested	11.1%	27 of 3,225 tested, few	25.0%	4 of 369 tested, few

EPA's method, non-detects at zero. SDWIS 2026 population and primary source; purchased water counted with its source type.


Three separate things shape EPA's national figure

The share of all U.S. water systems over a 2024 limit depends on how tested results are projected to untested systems, which small systems were drawn, and what the laboratories could quantify. Each moves it by a few percentage points.



Systems with complete sample sets, EPA's method. The 3-size projection is EPA's footnote-5 weighting (7.6% here; EPA, with its own size split, reports 7.8%). The 4-size by source-water projection post-stratifies by SDWIS population (500 or fewer, 501-3,299, 3,300-10,000, over 10,000) and primary source (GW, SW, GU), approximating EPA's stratified design; EPA's own strata and selection probabilities are not public. Sampling variance: stratified estimator with finite population correction (Cochran 1977).

89% of the sampling uncertainty comes from one stratum: 187 tested ground-water systems serving 500 people or fewer, standing in for 37,383 such systems, 56% of all the systems in the country.



Counted in people, sampling barely matters, because the systems serving most people were all tested: 17.2% (16.6-17.8%) of the 339.2 million people served by public water systems are served by one over a limit on the measured values. The laboratories' reporting level still moves it, to as much as 24.4%.

Questions UCMR 5 can and can't answer

● Answered

Was PFAS found in a tested system?

3,595 of 10,313 systems (34.9%).

Did any sample reach a 2024 limit?

1,763 systems. Non-detects play no part.

How high do the highest results go?

95th percentile 4.8 ppt PFOA, 5.0 ppt PFOS.

Do EPA regions differ?

Yes, but region explains little: η^2 of 0.084 at most (PFOA).

Does source water matter?

Surface-water samples detect PFAS more often (25.9% vs 16.3%), but are not more often over a limit.

🌀 As a range

Is a system's average over a limit?

1,136 yes, 8,501 no, 492 undecidable.

What is the average level?

PFOA somewhere from 0.52 to 4.26 ppt.

What share of all U.S. systems is over?

EPA says 7.8%. Sampling, projection and non-detects together allow 6.4% to 14.3%.

How precise is that?

8.4% (6.4–11.1%) by size and source; most of the uncertainty is the smallest ground-water systems.

🌀 Out of reach

What is a typical level?

The median is below the line for every PFAS, everywhere.

Is it under 4 ppt, or under 2 ppt?

For PFOA and PFOS the line is at 4 ppt.

How are my state's small systems doing?

A median of 12.5 tested per state.

What about untested systems?

56,144 systems, 34.0 million people.

Is it getting better or worse?

One year of samples per system.

Is my tap water safe?

Entry points, not taps; no exposure or health data.

Where does it come from?

Patterns can only be "consistent with" a source.

Is the system out of compliance?

Compliance averages start in 2029.

How the numbers were made

EPA's UCMR 5 file (all 1,991,977 results, IDs read as text) joined to SDWIS 2026 Q3. A non-detect lies between zero and the reporting level, so every average is recomputed with non-detects at zero, at half the reporting level, and just under it. A system is over a limit on the measured values if its average with zeros reaches the threshold, and undecidable if only the higher values would take it there. EPA's method: complete sample sets, one result per sample event, 4.05 ppt for PFOA and PFOS, 15 ppt for PFHxS, PFNA and HFPO-DA, Hazard Index 1.5 with two or more detected. With zeros the counts reproduce EPA's 1,136 exactly. National shares are post-stratified by size and source water, with stratified variances and finite population corrections (Cochran 1977); the 15 systems without a source code (none tested) count as ground water. Source-water intervals come from bootstrap resamples of whole systems. Distribution statistics follow Helsel (2005): Kaplan–Meier with one reporting level per PFAS; ROS and MLE only with at least 8 detects and no more than 80% non-detects, or 50 detects and a good lognormal fit. Scripts: `nd_sensitivity.py`, `design_sourcewater.py`, `km_state_region.py`, `epa_comparison.py`.

Sources

EPA, UCMR 5 Data Summary: July 2026 (EPA 815-S-26-007), Table 4. 40 CFR 141.903(f)(1)(iv) and 141.902(a)(5) (eCFR, October 2026). PFAS drinking-water rule, 89 FR 32532 (April 26, 2024). EPA, PFAS NPDWR Initial Monitoring Quick Reference Guide (December 2024). UCMR 5 final rule, 86 FR 73131 (December 27, 2021). EPA, UCMR 5 public meeting slides (April 2021): sample allocation. Cochran, W.G. (1977) Sampling Techniques, 3rd ed., Wiley. Lohr, S.L. (2019) Sampling: Design and Analysis, 2nd ed., CRC. Helsel, D.R. (2005) Nondetects and Data Analysis, Wiley. Helsel, D.R. (2006) Fabricating data, Chemosphere 65:2434–2439. Helsel and Lee (2006), NADA for R workshop examples. Wilson (1927), JASA 22:209–212. EPA SDWIS, 2026 Q3.

These are screening figures from occurrence monitoring, not compliance findings, and they do not identify any source or cause.