

113

the week's average on the national index; not one of 309 station readings was rated Good

The worst day was Sunday 19 July, at 172. Measured, the week ran +32 points on last year; the city's own share is +22. Every figure is read from the feeds, never retyped.

15 / 309

READINGS POOR OR WORSE
station readings this week; not one was rated Good.

+22

POINTS, THE CITY'S OWN
against the same days last year, weather removed.

3/7


DAYS, ONE STATION FIRST
Anand Vihar, dirtiest every single day.

The air record

Weekly deep edition · weekly to 24 Jul · prepared by Ajay Maken MP and Team

The seven days at a glance

The seven days, coloured by category



READING

Reading: each column is one day, coloured by the category the index fell in. The amber rule on each column marks the level the city's own pollution justifies once the weather is removed. The gap between that rule and the top of the column is therefore the weather's share of the day. The green line at 50 is the top of the clean-air band.

CAUTION

Caution: a day with no bulletin is drawn as an empty slot rather than dropped, because a shortened week would misrepresent it.

SOURCE

Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026_through_2026-07-24.csv [1].

THE POINT

Delhi did not reach the Good band on any day this week. Of 309 station readings, not one was rated Good.

The week, day by day, in both vintages

Day	Published that morning	Stations then	Recomputed now	Stations now	Category	Weather-corrected
Sat 18 Jul	171	45	171	45	Moderate	99
Sun 19 Jul	172	45	172	45	Moderate	100
Mon 20 Jul	134	44	134	44	Moderate	99
Tue 21 Jul	90	43	90	43	Satisfactory	100
Wed 22 Jul	68	45	68	45	Satisfactory	99
Thu 23 Jul	75	44	75	44	Satisfactory	100
Fri 24 Jul	78	43	78	43	Satisfactory	100

READING Reading: the week, one row per day, in BOTH vintages, with the station count behind each figure printed beside it.

CAUTION Caution, why two columns: the second column is the figure published that morning; the fourth is the same day recalculated now. This week the two agree on every day. Both columns appear in every edition regardless, so their absence can never be read as good news.

SOURCE Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026_through_2026-07-24.csv [1].

THE POINT

Where the two vintages differ, the figure published that morning stands. This report does not overwrite the record.

The weather, or the city's own pollution?

+32

Points against the same days last year, as measured, weather included.

+22

Points that survive once the weather is removed: the part the city's own pollution produced.

READING

Reading: a calm, still week makes the air look worse and a windy, wet week makes it look better. Neither is anyone's policy. The corrected figure is the accountability signal: it answers whether the city's own pollution actually rose or fell. Each day is compared with a fifteen-day window around the same date last year, and those windows overlap heavily. That is deliberate steadiness, not seven independent verdicts.

CAUTION

Caution: the weather-corrected figure is a model estimate, reported without a margin because the honest margin is still being built. No claim on these pages depends on it alone.


SOURCE

Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026_through_2026-07-24.csv [1].

THE POINT


The difference between +32 and +22 is the weather's contribution this week, in whichever direction it fell.

Which winds bring the dirtiest air: PM2.5



READING	Reading: a pollution rose is a compass built from the last 90 days. The LENGTH of a spoke is how often the wind came from that direction. The COLOUR bands are how dirty the air was on those days. Days with a dominant NW wind (22 days) carried the highest typical fine-particle level; E winds (19 days) brought the cleanest air.
CAUTION	Caution: only sectors with at least 5 days are ranked, so a one-day wind cannot be crowned. The dirty and clean sectors are the ends of this period's range across 4 ranked sectors; no significance test stands behind the ordering.
SOURCE	Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026_through_2026-07-24.csv [1].

Which winds bring the dirtiest air: PM10



READING

Reading: read length and colour together. A long pale spoke is a common but clean wind; a shorter deep-red spoke is a less frequent but dirtier one.

CAUTION

Caution: the dirtiest wind is the sector with the highest typical reading, which need not be the longest spoke on the rose.

SOURCE


Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026_through_2026-07-24.csv [1].

Does wind direction really matter?

Daily readings, grouped by the day's dominant wind

Delhi PM2.5 and PM10 by dominant wind direction (last 90 days)

Daily city-averages mapped to the day's predominant wind (24h to 10AM IST). The day count is printed under each sector, and the 4 sectors below 5 days are drawn unfilled because they are too thin to rank.



READING

Reading: which sectors tend to sit high and which sit low. A raw difference between sectors can mislead. The cleaner-looking winds are often simply the stronger winds, and strong winds clear pollution wherever they come from.

CAUTION

Caution: the honest test is on the next page. It asks whether direction still explains pollution after wind speed and the depth of the atmosphere are already accounted for.

SOURCE

Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026-07-24.csv [1].

Does wind direction survive the honest test?

Month	Pollutant	Days tested	Sectors	Chance likelihood	After the screen	Direction matters?
April	PM2.5	30	7	0.467	0.622	No
April	PM10	30	7	0.575	0.658	No
May	PM2.5	31	5	0.081	0.162	No
May	PM10	31	5	0.001	0.005	Yes
June	PM2.5	30	7	0.001	0.005	Yes
June	PM10	30	7	0.193	0.309	No
July	PM2.5	24	7	0.806	0.806	No
July	PM10	24	7	0.017	0.045	Yes

READING

Reading: whether wind direction still explains pollution after wind speed and the depth of the atmosphere are accounted for. A lower figure means direction is doing more work, and the screen makes the test harder to pass.

CAUTION

Caution: a persistence-corrected version of this test was built and tested. At a single month's sample size it proved unreliable in both directions. This table therefore reports the ordinary test with its sample size shown; the corrected version is held back until it can be made dependable.

SOURCE

Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026_through_2026-07-24.csv [1].

THE POINT

Only if direction survives this adjustment can a sector be said to bring pollution in. Otherwise it merely labels a calm or a windy day.

Weather against PM2.5, month by month

Against the same days last year

	January	February	March	April	May	June	July
Wind speed <i>cleaner air when higher</i>	-0.08 (-0.06)	-0.23 (-0.28)	-0.26 (-0.28)	-0.18 (-0.15)	-0.18 (-0.18)	-0.21 (-0.18)	-0.40 (-0.43)
Temperature <i>no fixed direction</i>	+0.43 (+0.24)	+0.45 (+0.47)	+0.70 (+0.75)	+0.65 (+0.45)	+0.66 (+0.62)	+0.70 (+0.68)	+0.68 (+0.66)
Depth of the mixing layer <i>cleaner air when higher</i>	-0.31 (-0.22)	+0.40 (+0.34)	+0.39 (+0.38)	+0.12 (+0.14)	+0.30 (+0.25)	+0.43 (+0.45)	-0.24 (-0.27)
Rainfall <i>cleaner air when higher</i>	+0.08 (+0.18)	-0.27 (-0.28)	-0.15 (-0.31)	-0.45 (-0.32)	-0.16 (-0.29)	-0.27 (-0.31)	-0.40 (-0.52)
Morning inversion strength <i>cleaner air when lower</i>	-0.37 (-0.19)	+0.07 (+0.03)	-0.03 (-0.00)	+0.22 (+0.37)	+0.06 (+0.18)	-0.09 (+0.00)	not measured
Morning inversion height <i>cleaner air when higher</i>	+0.08 (+0.07)	-0.54 (-0.16)	-0.00 (-0.02)	-0.45 (-0.38)	-0.36 (-0.48)	-0.33 (-0.40)	not measured

Against the day before

	January	February	March	April	May	June	July
Wind speed <i>cleaner air when higher</i>	-0.63 (-0.67)	-0.46 (-0.41)	-0.60 (-0.57)	-0.33 (-0.34)	-0.46 (-0.40)	-0.56 (-0.53)	+0.11 (+0.04)
Temperature <i>no fixed direction</i>	+0.18 (+0.15)	+0.00 (+0.00)	+0.18 (+0.00)	+0.38 (+0.46)	+0.32 (+0.32)	+0.21 (+0.22)	+0.56 (+0.66)
Depth of the mixing layer <i>cleaner air when higher</i>	-0.53 (-0.56)	-0.35 (-0.24)	-0.36 (-0.38)	-0.05 (-0.03)	-0.18 (-0.19)	-0.08 (-0.10)	+0.15 (+0.18)
Rainfall <i>cleaner air when higher</i>	-0.01 (-0.06)	+0.35 (-0.09)	-0.15 (-0.44)	-0.04 (-0.21)	-0.32 (-0.28)	-0.39 (-0.26)	-0.20 (-0.56)
Morning inversion strength <i>cleaner air when lower</i>	-0.11 (+0.08)	+0.40 (+0.04)	+0.25 (+0.30)	+0.22 (+0.17)	-0.01 (+0.02)	-0.16 (-0.09)	not measured
Morning inversion height <i>cleaner air when higher</i>	-0.09 (-0.19)	-0.26 (+0.10)	-0.29 (-0.10)	-0.15 (-0.16)	-0.27 (-0.18)	-0.14 (+0.29)	not measured

READING

Reading: each square asks whether, when the weather ran unusual, the pollution ran unusual in step. A coloured square is a link this report will stand behind. A white square with a dotted edge failed the reliability screen; its number prints for completeness only. The second line under each weather name says which direction usually helps Delhi's air.

CAUTION

Caution: a number in brackets is a second measure of the same link, one not thrown by a handful of extreme days. Where the two differ materially, the bracketed figure is the more trustworthy.

SOURCE

Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026-07-24.csv [1].

Weather against PM10, month by month

Against the same days last year

	January	February	March	April	May	June	July
Wind speed <i>cleaner air when higher</i>	-0.08 (-0.04)	-0.24 (-0.22)	-0.01 (-0.06)	-0.06 (-0.04)	-0.21 (-0.19)	-0.15 (-0.02)	-0.22 (-0.35)
Temperature <i>no fixed direction</i>	+0.39 (+0.18)	+0.58 (+0.62)	+0.67 (+0.70)	+0.74 (+0.47)	+0.78 (+0.71)	+0.76 (+0.76)	+0.82 (+0.87)
Depth of the mixing layer <i>cleaner air when higher</i>	-0.35 (-0.30)	+0.34 (+0.28)	+0.49 (+0.47)	+0.03 (+0.01)	+0.34 (+0.28)	+0.51 (+0.52)	+0.01 (-0.03)
Rainfall <i>cleaner air when higher</i>	-0.03 (-0.02)	-0.27 (-0.25)	-0.26 (-0.39)	-0.40 (-0.20)	-0.06 (-0.23)	-0.32 (-0.44)	-0.46 (-0.72)
Morning inversion strength <i>cleaner air when lower</i>	-0.28 (-0.04)	+0.10 (+0.08)	-0.05 (-0.00)	+0.12 (+0.29)	+0.08 (+0.06)	-0.01 (+0.10)	not measured
Morning inversion height <i>cleaner air when higher</i>	+0.04 (-0.01)	-0.48 (-0.08)	+0.02 (+0.04)	-0.27 (-0.16)	-0.25 (-0.41)	-0.29 (-0.22)	not measured

Against the day before

	January	February	March	April	May	June	July
Wind speed <i>cleaner air when higher</i>	-0.73 (-0.77)	-0.55 (-0.59)	-0.65 (-0.64)	-0.11 (-0.20)	-0.24 (-0.16)	-0.27 (-0.16)	+0.25 (+0.09)
Temperature <i>no fixed direction</i>	+0.20 (+0.11)	+0.19 (+0.10)	+0.31 (+0.16)	+0.45 (+0.57)	+0.39 (+0.41)	+0.51 (+0.36)	+0.60 (+0.73)
Depth of the mixing layer <i>cleaner air when higher</i>	-0.60 (-0.59)	-0.39 (-0.31)	-0.45 (-0.42)	-0.01 (-0.04)	-0.31 (-0.22)	+0.23 (+0.03)	+0.36 (+0.25)
Rainfall <i>cleaner air when higher</i>	-0.22 (-0.12)	+0.17 (-0.12)	-0.23 (-0.45)	-0.08 (+0.01)	-0.29 (-0.22)	-0.32 (-0.17)	-0.17 (-0.63)
Morning inversion strength <i>cleaner air when lower</i>	-0.05 (+0.20)	+0.30 (+0.01)	+0.28 (+0.29)	+0.11 (+0.01)	-0.03 (-0.01)	-0.16 (-0.24)	not measured
Morning inversion height <i>cleaner air when higher</i>	-0.07 (-0.08)	-0.16 (+0.18)	-0.45 (-0.21)	+0.02 (+0.07)	-0.06 (+0.07)	-0.15 (+0.29)	not measured

READING

Reading: each square asks whether, when the weather ran unusual, the pollution ran unusual in step. A coloured square is a link this report will stand behind. A white square with a dotted edge failed the reliability screen; its number prints for completeness only. The second line under each weather name says which direction usually helps Delhi's air.

CAUTION

Caution: a number in brackets is a second measure of the same link, one not thrown by a handful of extreme days. Where the two differ materially, the bracketed figure is the more trustworthy.

SOURCE

Weekly export weekly_export_2026-07-24.json from build_weekly_deep.R; the week's figures from the daily narration feeds; input CREA_Delhi_Raw_Long_2026-07-24.csv [1].

How to read these grids

What the grids measure

Neither grid is about how dirty the air was. Both are about departures: how far a day sat above or below what would be normal for it. A month can be filthy throughout and still show no link here, if its filth did not rise and fall with the weather.

The two grids can disagree. When they do, it is usually because they answer different questions, not because one of them is wrong. Read each for the question it asks.

One square, read in full

Take the square where wind speed meets fine particles in June. Suppose it reads about minus one half. Here is what produces a number like that.

On one day the wind drops well below normal and the pollution climbs above normal. That is one day, and one day settles nothing. On another the wind picks up and the pollution falls. Now take every consecutive pair of days in the month. On most of them the wind and the pollution moved in opposite directions, the direction Delhi would want. On the rest they moved the same way.

The single number summarises all of those days at once. The minus sign says that when the wind picked up, the pollution usually fell. The size says how firmly. Squared, it becomes the plainest statement of strength. A correlation of about minus

one half means the two shared roughly thirty per cent of their day-to-day movement. The other seventy per cent moved independently.

That seventy per cent is not a rounding error. There are days in every month when the wind rises and the pollution rises anyway. A correlation is a tendency across a month. It is never a rule that must hold on a given morning.

Three things a square does not say

It does not say the weather caused the change. In Delhi the wind, the rain, the temperature and the depth of the mixing layer all rise and fall together. Any one of them will therefore appear to move with pollution partly because the others did.

It does not say the weather is why Delhi's air is bad. Wind decides how much of a day's pollution drifts away and how much settles over the city. It does not decide how much the city puts into the air to begin with. Take a month when the wind had been generous on every single day. Delhi's vehicles, construction sites, industry and waste fires would still have emitted precisely what they emitted.

And it does not carry to another month. The same pairing can be strong in one month and absent in the next. That is exactly why this report prints every month rather than a single figure for the year.

The two tests, and why there are two

Why one significance test is not enough

Delhi's weather does not start afresh each morning. A still spell lasts three or four days, and so does a windy one, and the pollution beneath it behaves the same way. Thirty days of weather are therefore not thirty separate pieces of evidence. The ordinary significance test assumes that they are, and so it can credit a month with more confidence than the month has actually earned.

This report therefore runs a second test alongside the ordinary one, a correction published by Whitney Newey and Kenneth West. Rather than treating each day as fresh, it measures how much each day still resembles the day before it, the day two before and the day three before. It then rebuilds the margin around the correlation to match. The number of days it looks back is fixed by a rule and not chosen by hand, and for a month of about thirty days that rule gives three.

The correction does not always make this report stricter, and we will not pretend that it does. On the day-on-day grid each reading is already a change from the day before, and that subtraction has usually over-corrected for the carry-over. When that happens the rebuilt margin comes out narrower than the ordinary one, and the correction makes a finding easier to reach rather than harder. Measured across all 160 tests in this edition, the correction moved the answer towards caution in only 50 of them.

A correction that can loosen the standard cannot be the only guard. So this report computes both answers and keeps whichever of the two is the more cautious. A square has to satisfy both tests before it is allowed to carry colour. That is a stricter standard than either test on its own, and it cannot be embarrassed by either one misbehaving.

Why a grid needs a second screen

One grid on these pages is not one test. For most months it is twelve, being six weather measurements set against two pollutants. Across the months and both grids, this edition runs 160 of these tests in all.

That is the difficulty. Toss a coin often enough and some run of heads will look remarkable. Compare enough pairs of unrelated numbers and some pair will look connected. Testing twelve things at once and then pointing at the most striking is how a finding gets manufactured that was never there.

The screen applied here was published by Yoav Benjamini and Yosef Hochberg. It ranks a grid's results from the most striking to the least and raises the bar for each in turn. A result must be considerably more striking to be believed when it is one of twelve than when it stands alone. The bar is set at ten per cent. The plain meaning is this: among the squares this report allows to carry colour, about one in ten will be a coincidence rather than a real link. We cannot tell you which one. Each month and each grid is screened on its own, so a weak month cannot borrow credibility from a strong one.

A square computed on fewer than 20 usable pairs of days is not tested at all. A margin that corrects for carry-over cannot be trusted on a handful of days. Its number still prints, in grey. Of the 160 squares in this edition, 33 carry colour.

One honest limit. This screen works cleanly when the tests on a grid do not conspire with one another. Delhi's weather does not fully oblige. Wind, rain, temperature and mixing depth rise and fall together, so the twelve tests on a grid are related to each other. Read the ten per cent as a sound working guide rather than as a guarantee.

Method, in plain English

The daily building block

Every pollution number is a 24-hour average ending at 10 in the morning, the same window Delhi's pollution control board uses. A day in this report therefore means yesterday morning to this morning. A monitoring station counts on a given day only if it reported for at least sixteen of those twenty-four hours. That keeps a half-broken sensor from skewing the city figure.

Why only two pollutants

Delhi's monitors also measure four gases. That feed froze on 23 June 2026 and its numbers cannot be trusted, so this edition reports fine particles and coarse dust only. Their measurements are independent of that fault. This is also why the daily report shows blank days in late June where this one does not. A national index is a composite that needs the gas readings. It could not be formed on days when the particle monitors were reporting perfectly well.

The week's figures

The week is the average of the figures published each morning, never a fresh calculation of the whole week at once. Those daily figures were produced by one version of the weather-correction model. Recomputing the week now would use a model none of those mornings used.

What the correction can and cannot say

The weather-corrected figure is an estimate from a model, not a measurement. It is reported without a margin in this edition. The honest way to measure that margin is still being built, and quoting a made-up one would be worse than quoting none. No claim on these pages depends on it. The week's verdict rests on a gap of many points. The widespread finding is a count of stations, and the persistence finding is a count of days.

The screen, in one paragraph

Each square is tested twice: once by an ordinary significance test, and once by a test that tolerates the way Delhi's weather persists for days at a time. The more cautious of the two answers counts. Every grid is then screened again for the fact that it runs twelve tests at once, at a ten per cent false-discovery rate. A square computed on fewer than 20 usable pairs of days is not tested at all. Squares that fail print their number in grey.

Provenance

Pollution figures come from the CREA compilation of the pollution control board's station data. Weather comes from the ERA5 and Open-Meteo hourly record for Delhi, plus the Safdarjung weather balloon. Every figure in this report was recomputed from those files at the moment it was built. The exact input file is named on every page.

Full provenance: input CREA_Delhi_Raw_Long_2026_through_2026-07-24.csv · sha256 6ef9d52b95e9 · 1,596,498 rows