

Change Tracker MONTHLY

The monthly momentum register that tells you what is materially shifting on your topic, trend by trend, cycle over cycle, mapped onto an interactive Trend Radar and read across a Cross-impacts grid that shows which trends are driving which.

Signal Scanner → **Change Tracker** → Decision Intelligence → Action Triggers

1 What it is

The Change Tracker is the assessment layer of a Shaping Tomorrow continuous-intelligence programme. Where the weekly [Signal Scanner](#) detects individual weak signals, the Change Tracker steps back monthly and grades the whole register: which of your topic's tracked trends are accelerating, which are stable, which are fading, what is new this cycle, and what pattern the movements make when read together. Every edition maps that register onto an integrated, interactive **Trend Radar**, the same assessments drawn as a horizon chart a board can read in one glance, and carries a **Cross-impacts** assessment: a sourced grid of which trends are speeding up or slowing down which others, resolved into the three collisions a leadership team most needs to see together.

The Shaping Tomorrow intelligence cycle

What changed this week, what is moving this month, what to decide this quarter, and one page that always shows where things stand. This deliverable's lane is highlighted.



Each tracker re-grades the same register against the month's scans. Its movements feed the Live Summary as they happen, and the closing Decision Intelligence briefing and Action Triggers at the end of the cycle.


A recent edition, *"Agents escape the sandbox as Europe enforces and Washington reviews in private"*, tracks ten trends across Artificial Intelligence and Automation: each graded on heat, horizon and momentum in the register table, plotted on the radar with three wild cards on the outer rim, linked to one another across fifteen sourced cross-impacts, then read into three pattern themes and a decision-facing implications verdict.

THE PROBLEM IT SOLVES

Detection without assessment produces noise: a stream of interesting signals with no sense of which are gathering force and which are quietly dying. The Change Tracker is the discipline that turns a watchlist into a managed register: every trend re-graded against fresh evidence each cycle, direction changes called explicitly, new entries badged, departures marked, and the whole diff drawn onto the radar as drift tails so you always see what moved, what didn't, and what just arrived or left. Grading one trend at a time still leaves one question open: which of them is driving the others. The Cross-impacts grid answers it from the same evidence, so a loop between two trends is on the page before it is in the news.

The Trend Radar, from a live edition

The Health, Life Sciences & Care Systems register, 11 September 2026, exactly as published. Rings are time (Now · Next · Later), sectors are the cycle's pattern themes, blip size is heat, arrows are momentum, dashed ghosts show last cycle's position, and the diamonds on the dashed rim are wild cards. On the live page, hovering any blip opens its full register entry.



2 Why clients choose it

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
A managed register, not a feed	The same trends re-graded every cycle, with direction changes, new entries and departures called explicitly and drawn as drift on the radar.	You always know the current state of your topic, and can see, spatially, when the programme called each turn.
Momentum, the missing dimension	Most intelligence tells you what exists; the tracker shows you what is speeding up, slowing down, arriving and leaving, with arrows pointing at your planning horizon.	Timing is where strategic value lives. Knowing a constraint is real matters less than knowing it just accelerated.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
The loops, not just the list	The Cross-impacts grid shows which trends are driving which, with every link sourced, names the trend whose shifts spread furthest, and sets out the three collisions that matter most in plain English.	A register of ten independent trends is a fiction; the real risk is usually two of them feeding each other. The grid is where that becomes visible before it becomes a headline.
Both tails covered	Trends in the register, discontinuities as diamonds on the radar rim, each with a dated likelihood and a pathway showing how far it has already travelled.	The register that watches what it tracks and what could ambush it is the one that prevents the “nobody saw it coming” meeting.
The programme's connective tissue	Trackers grade what the Signal Scanner detected, including each scan's dated call once it falls due, feed the Live Summary as they publish, and build the movement record that Decision Intelligence and Action Triggers synthesise.	Each monthly edition compounds the programme's memory: the asset a subscription builds and a one-off report never can.

FROM A RECENT EDITION (3 SEPTEMBER 2026)

In one monthly read, the tracker: graded ten trends on heat, horizon and momentum with a linked source in every row, six of them accelerating; mapped them onto the Trend Radar with three wild cards on the rim; found fifteen sourced cross-impacts between them, nine that speed another trend up and six that slow one down, four pairs working both ways; named the three collisions that matter most, including two lab containment breaches turning agent monitoring into a supervisory requirement, and capital spending and a valuation correction feeding each other as one loop; resolved the register into three pattern themes and the cycle's headline; and closed with a decision-facing implications verdict, on sixteen tier-tagged sources.

3 The register discipline

The Change Tracker's power is that it is the same register, revisited. Each edition carries the previous cycle's title in its continuity line, re-grades every trend against what actually happened, badges new entrants, marks departures, and names the cycle's single biggest movement in its

headline. The Trend Radar is that comparison's reader-facing surface: a dashed ghost blip shows each changed trend's previous state, a dotted tail leads to where it sits now, and a grey crossed ghost marks a trend carried last cycle but absent from this one. The cross-impact links are carried the same way, so the next edition can say which links strengthened, weakened or dropped. From October 2026 each tracker also grades the Signal Scanner calls that have come due, on that cycle's sourced evidence only, as confirmed, on track, partly, unresolved or broken, and the grades are copied unedited to the programme's public Track record. Over a year, the trackers accumulate into a documented history of how your topic moved, how its parts acted on each other, and when the programme called each turn: the evidence base the quarterly [Decision Intelligence](#) briefings and end-of-cycle [Action Triggers](#) then draw on.

4 The design system: why the format itself is a benefit

The tracker is built to be read radar-first and audited table-second, and its design carries that:

- **Geometry as grading:** rings, sectors, blip size, arrows, tails and rim diamonds each encode one register field, in the same polarity palette as our Futures Wheel, so the whole family of deliverables reads as one system.
- **Soft grading grammar:** momentum pills, heat chips and cluster chips use tinted fills with coloured text rather than solid blocks, in identical form across the pill row, the table and the hover panel.
- **The panel teaches first:** the radar's side panel defaults to a how-to-read guide, swaps to full register entries on hover, and brings the guide back with a ? button, so the chart needs no separate manual. The Cross-impacts grid follows the same rule: its explanation sits above it in one plain sentence, and every cell explains itself on hover.
- **One chart, one grid:** the Trend Radar is the tracker's only chart. The Cross-impacts grid is a generated table in the tracker's own palette, red for speeds up and green for slows down, so it introduces no new visual language, prints as it reads, and is fully accessible.
- **Print-disciplined:** in print the interactive radar yields to the momentum table, which carries the complete register including wild cards and departures, and the grid and its three collisions print in full; the PDF edition loses nothing.

5 Anatomy of the deliverable: element by element

The background detail for readers who want to see exactly what is in each edition and why, section by section.

[Click a section to expand it.](#)

The header: headline, standfirst and continuity

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Two-clause headline	The cycle's verdict in one line: what hardened, what advanced, what cooled.	The tracker leads with its conclusion, not its contents. You know the cycle's biggest movements before scrolling.
Standfirst	One italic sentence compressing the whole edition, with its reasoning attached.	The 15-second version for the executive who will read the radar and nothing else.
How this was made	A line under the title, and a card at the foot of the tracker, linking to Shaping Tomorrow's six-stage method: Scan, Filter, Stress-test, Judge, Trigger and Actionable Intelligence. The card also links to how each edition is verified and to the public Track record of graded calls.	A reader who meets one edition on its own can see how it was produced, and how earlier calls turned out, before deciding how far to rely on it.
Series-continuity line	"Edition 3 · previous cycle: 'Bank clearing layer for tokenised deposits arrives as the digital euro nears a vote'".	Every edition situates itself in the series. You can see, from the two titles alone, how the story moved in one month.
Consensus-and-change opening	A short paragraph granting what still holds, then naming precisely what changed, with sources inline.	Same intellectual manners as the Signal Scanner: no straw men, no re-litigating the settled, straight to the delta.

The Trend Radar: the register on the horizon

The signature visual, in the same design language as our Futures Wheel. Every element of the chart is data:

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Geometry is meaning	Rings = when a trend bites (Now · Next · Later; closer to the centre means sooner). Sectors = the cycle's pattern themes, in the Futures Wheel palette. Blip size and stroke = heat (impact × likelihood).	The whole register in one glance: what is close, what is hot, and which theme it belongs to, before a word is read.
Momentum arrows	An arrow pointing toward the centre means Accelerating; outward means Fading; no arrow means Stable.	Momentum becomes geometry: the trends closing in on you literally point at the centre of your planning horizon.
Drift tails and ghosts	A dashed ghost blip (sized to last cycle's heat) with a dotted tail marks every trend whose grades changed; a grey crossed ghost marks one that left the register. Trends that held their grades get no tail.	Cycle-over-cycle honesty, drawn on the chart. Your eye goes straight to what moved, and every marker traces to the prior edition's literal record.
Register badges	A NEW badge on entrants; an orange ↘ on trends that cooled a stage; the momentum pill row above the radar (6 Accelerating · 2 Stable · 1 Fading, plus the cycle's flags) in soft tinted chips.	The three questions a returning reader has (what arrived, what cooled, what is the temperature) answered before the first hover.
Hover panel with guide	The side panel opens on a “How to read this radar” card; hovering any blip replaces it with the trend's full register entry. From October 2026 that entry leads with “What moved this cycle”, then the grade changes since the prior cycle and the trend's likelihood band with its percentage; a ? button recalls the guide.	Self-teaching for first-time readers, instant for returning ones, and fully usable on touch. Nothing needs a manual.
Expand view	A full-screen reading mode with the enlarged radar and a docked detail panel; hover and pin work identically there.	Made for the board-room screen share: the radar becomes the meeting's working surface, not just a page graphic.

One row per tracked trend, six disciplined columns; the radar’s static, print and accessibility twin:

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Six-column register	Trend Cluster Horizon Heat Momentum Evidence, with soft tinted chips for cluster, heat and momentum.	Instant triage down the page: the same grading grammar as the radar, in scannable rows a diligence team can audit cell by cell.
What moved this cycle NEW	From the October 2026 cycle, a one-line answer sits in grey under every trend name, with an arrow coloured by momentum: “Nine more deals on 31 August take coverage to 26 companies and 89% of the branded market”, or, just as plainly, “No movement: four months after the provisional deal, still no formal vote”. The Evidence column keeps the standing position, so the two never repeat each other.	The question a returning reader actually has, “what changed since last month?”, is answered row by row without opening a single source. A held trend says so honestly instead of hiding behind restated evidence.
Likelihood under the heat chip NEW	Each heat chip now carries the likelihood behind it, with its percentage band: “Very likely, over 80%”. The same bands the Signal Scanner uses, so every grade is scorable later.	You can see why a trend is rated hot, and challenge it: heat stops being a black box and becomes impact times a stated, checkable probability.
Register markers in the rows	The NEW chip and the ↘ cooled marker sit on the trend names; wild cards and any trend that left the register appear as their own rows.	The table is the complete register, tails included: nothing on the radar is missing from the rows, and vice versa.
Evidence with source inline	Each row’s evidence closes in Change Tracker citation form: the linked publisher and month in parentheses.	Every claim is one click from its source, in the same citation grammar as the tracker’s prose. Due-diligence-proof at a glance.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Radar-table consistency	Radar and table are generated from one register: same heat bands, same momentum, same horizons.	The chart can never disagree with the rows. In print, the table carries the full register so the PDF loses nothing.

The pattern narrative: reading across the rows

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Typed theme cards	The trends read together into colour-coded theme cards: Transformation Driver, Emerging Opportunity, Evolving Risk, each with sources inline.	This is the analysis a chart cannot do: the insight that three trends are one story. It is also where the tracker earns its headline, and it names the radar's sectors.
Signals gaining momentum	A ranked list of the cycle's risers, each with its concrete escalation consequence.	Not just "this is rising" but "here is what forces a decision if it keeps rising": the sentence that moves a trend onto an agenda.

Cross-impacts: how the trends act on each other

NEW, SEPTEMBER 2026

Each trend is graded one at a time. The Cross-impacts block is where the Change Tracker asks what one trend does to another, scored after the register and the pattern narrative are settled, from the same evidence base, and built to be read by someone who has never seen the grid before:

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
The grid	One row and one column per tracked trend. Read along a row to see what that trend does to the others: a red cell means it speeds the other up, a green cell means it slows it down, and the number (1 to 3) says how strongly. The diagonal is hatched; blank means no evidence of a link.	The register graded one trend at a time cannot show you a loop. The grid does, in one read, and it is honest about what is not connected: most cells in a ten-trend register stay blank, and that is the finding.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Direction and strength	Every link is directed, so a pair that acts both ways appears as two mirrored cells and is marked “works both ways”. Strength follows a fixed rubric: 3 where the second trend’s momentum this cycle is largely explained by the first, 2 for a clear sourced contribution, 1 for an evidenced nudge.	Strength is defined, not felt. A 3 has a dated causal chain behind it; a 1 is a plausible nudge and says so.
Blank means unevidenced	A cell is filled only when a mechanism sentence traces to at least one source in the cycle’s bibliography. Six to fifteen links per cycle; fewer means the pass was skipped, more means the grid is being filled to look busy.	No source, no cell. The grid inherits the tracker’s sourcing discipline rather than becoming the one place where opinion creeps in.
The standout line	One plain sentence under the link counts naming the trend that drives the most change in the others and the trend that receives the most.	Tells you where a shift spreads furthest and which trend moves when the others do, without asking you to add up the grid.
Hover panel with mechanism and sources	Hovering or tapping a coloured cell opens the reason: the mechanism sentence and the sources behind it. A chip row above the grid counts the cycle’s links: how many speed up, how many slow down, how many pairs work both ways.	Every number in the grid is one hover from its evidence, in the same citation grammar as the rest of the tracker.
Three collisions	Beneath the grid, exactly three collision blocks: the links a reader most needs to see together, each with a plain-English title (“Capex growth and a valuation correction feed each other”), a read of at most 70 words, and a Watch line naming the observable that would confirm or break the link.	The grid is the evidence; the collisions are the argument. Each title is written to be understood by a trustee who has not looked at the grid, and each Watch line is a tripwire the next cycle checks.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Links carried cycle to cycle	The link list is stored with the edition's register, so the next tracker can say which links strengthened, weakened or dropped.	Interactions get the same cycle-over-cycle discipline as the trends themselves. A loop that is tightening is visible as it tightens.

Wild Cards to Watch: the discontinuity layer

UPDATED, OCTOBER 2026

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Wild-card cards	Low-probability, high-impact discontinuities, typed Disruptive Risk or Disruptive Opportunity with an impact rating, and plotted as dashed diamonds on the radar's outer rim.	The tracker covers the tails, not just the trends, and the radar keeps them visibly on the edge of the picture where they belong.
A dated likelihood NEW	From October 2026 every card states "Likelihood: 5 to 30% by 2028": a wild card is by definition 5 to 30% likely within its window, above that it joins the tracked trends, below it leaves the page.	The tail risks carry odds and a deadline, so they too can be graded later rather than lingering unfalsifiably.
Surprise characteristics	What would make the event a genuine surprise rather than a trend extension.	Forces precision about the discontinuity itself: the analytic step most wild-card lists skip.
The pathway, and where we are NEW	From October 2026 each card's early-warning indicators are set out as an ordered chain of four to six steps from today to the wild card landing, the last in bold beginning "The wild card lands:". A dashed WE ARE HERE marker sits after the last step that has already happened, and every step above it carries a "Seen" chip naming the sourced fact that shows it, such as "Seen: up 6.7% in 2025".	You see not just what to watch but how far along the road a tail risk already is. The next tracker moves the marker rather than redrawing the chain, so its progress becomes a graded claim, like a drift tail on the radar.

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Evidence paragraph	Each card grounded in sourced fact, with the tracker’s inline citation form.	Wild cards argued from data, not imagination. The section reads as analysis, not brainstorming.

The close: implications and evidence base

ELEMENT	WHAT THE READER SEES	WHY IT MATTERS TO YOU
Implications box	A single decision-facing paragraph: what to prepare, what to pilot, what to monitor, through named decision points.	The tracker ends by telling you where to point your attention next month, which is precisely what a monitoring product owes you.
Tiered bibliography	Every source tier-tagged (T1 governments/regulators → T4 vendor, “directional corroboration only”), dated, and linked.	The same four-tier taxonomy as every Shaping Tomorrow deliverable: monthly cadence, due-diligence sourcing.

Change Trackers publish monthly within each cycle of a Single-Topic or Multi-Topic continuous-intelligence programme, assessing the signals the weekly [Signal Scanner](#) detects and feeding the [Decision Intelligence](#) briefings and [Action Triggers](#) syntheses that decide and act on them, as well as the programme’s Live Summary, which is refreshed after every deliverable. Every edition ships with the integrated Trend Radar and the Cross-impacts assessment as standard. The tracked register, cadence and distribution are tailored to your topics and planning calendar.

Discuss a Tailored Programme

Request a Sample