

The analytics capability layer

Four progressive capabilities sit in the analytics and BI layer of the maturity model, and each answers a different business question: what happened, why it happened, what will happen, what we should do about it. Each also builds directly on the governed data foundation and semantic layer below it.

This is a working reference, not a summary. The targets below are illustrative starting points; the questions are the useful part. Take the one capability you are actively investing in, answer its questions honestly, and fill in its current-state column. That filled-in gap, not the illustrative target, is what should drive your next prioritization conversation.

THE FOUR QUESTIONS BEHIND EVERY ROW

What is it? The outcome the enterprise can produce or protect, not the tools used to produce it.

How well must we do it? The performance expectation, and the gap between current and desired state.

Why does it matter? Which lever it pulls: revenue, cost, service, compliance, or future positioning.

What could go wrong? In performing it, or in failing to perform it at all.

The four capabilities at a glance

CAPABILITY	THE OUTCOME IT PRODUCES	WHERE ITS VALUE CONCENTRATES	THE QUESTION THAT EXPOSES YOUR GAP
1 · Descriptive Analytics	One consistent, agreed-upon picture of the business, delivered on a predictable, trusted cadence.	Service and cost. The foundation everything else inherits.	If the CFO and a regional manager pulled last quarter's revenue right now, would they match?
2 · Diagnostic Analytics	When a number moves, the region, product or driver behind it found in minutes rather than by filing a request.	Analyst capacity, and speed to corrective action.	How many of last month's ad hoc requests were drill-downs a business user could have run?
3 · Predictive Analytics	Forecast values and likelihoods from historical patterns, embedded back into the reports people already trust.	Revenue and risk, ahead of the event.	If this month's forecast missed by 30%, would anyone know before next quarter's planning?
4 · Prescriptive Optimization	The best available action recommended or executed against a forecast, constraints and an objective, with the outcome fed back.	Direct revenue and cost, at the highest stakes.	Of last quarter's recommendations, how many were executed, and how many were confirmed to have worked?

CAPABILITY 1

Descriptive Analytics

The ability to summarize what has happened and what is happening now (trends, totals, KPIs, dashboards built on governed, conformed data) and to deliver that view to the people who need it on a predictable cadence. Whether it renders in one BI tool or another is an implementation choice below the capability line.

This is the entry point for nearly every analytics program, and the capability every other one in the layer depends on for a shared, certified starting point.

Improve service	Replaces conflicting spreadsheets with one trusted view, so nobody re-litigates whose numbers are right.
Reduce cost	Removes the recurring labor of assembling and reconciling reports across teams every cycle.
Achieve compliance	Certified, governed dashboards create an auditable record of what the business knew and when.
Position for future	Diagnostic, predictive and prescriptive work is only as trustworthy as the descriptive foundation under it.

PERFORMANCE DIMENSION	ILLUSTRATIVE TARGET	OUR CURRENT STATE
Refresh reliability	Scheduled refresh succeeds on $\geq 99\%$ of runs	
Time to a trusted number	A new certified metric or dashboard ships within an agreed cycle time of request	
Single-source-of-truth coverage	Share of key metrics with one certified definition, against competing shadow versions	
Adoption	Share of the target audience using it on a recurring basis, not just at launch	

ASK BEFORE YOU SET A TARGET

- How many versions of the truth exist today for our top metrics?
- If the CFO and a regional manager pulled last quarter's revenue right now, would they match? If not, the gap isn't dashboard design; it is the absence of a certified semantic layer, and no amount of visualization polish will fix it.
- Are reports built on a modeled semantic layer, or directly on raw source tables that break downstream with every schema change?

CAPABILITY 2

Diagnostic Analytics

The ability to explain why a metric moved: drill-down, decomposition, cohorting and root-cause analysis performed on the same governed model that produces the descriptive view, without an analyst having to intervene for every question.

This capability introduces no new data. It adds a layer of explanatory power on top of data that is already trusted.

Reduce cost	Deflects analyst backlog by letting business users self-serve the root-cause questions that would otherwise become tickets.
Improve service	Shrinks the distance between “something changed” and “we know why” from days of back-and-forth to a few clicks.
Protect revenue	Faster root-cause work on a revenue or margin drop means corrective action lands while it still matters.
Position for future	The segmentation, anomaly-detection and time-intelligence logic built here is what Predictive needs for features and baselines.

PERFORMANCE DIMENSION	ILLUSTRATIVE TARGET	OUR CURRENT STATE
Self-service root-cause rate	Share of “why did this change” questions resolved without analyst involvement	
Time to explanation	Median time from an observed anomaly to an identified driver	
Anomaly detection coverage	Share of key metrics with automated outlier detection rather than visual inspection	
Ad hoc request volume	Trending down as self-service diagnostic coverage increases	

ASK BEFORE YOU SET A TARGET

- Of last month's ad hoc analysis requests, how many were a drill-down a business user could have done themselves? A high number means the gap is capability, not headcount.
- Would a real anomaly (a fraud spike, a supply disruption, a churn wave) be caught if nobody happened to open the right chart that week?
- Is diagnostic logic centralized in the model, or duplicated across reports so two root-cause answers can disagree?

CAPABILITY 3

Predictive Analytics

The ability to forecast future values and estimate likelihoods from historical patterns: trend projection, propensity scoring, demand and risk forecasting, anywhere from lightweight in-tool forecasts to fully productionized, monitored pipelines.

This is where the architecture starts to diverge from BI tooling alone and begins depending on feature engineering, training and monitoring practices shared with the ML layer above.

Protect revenue	Demand forecasting lets the business plan inventory, staffing and pricing ahead of need rather than after a shortfall.
Reduce cost	Churn and risk propensity scoring allow intervention before a loss instead of paying its full cost afterward.
Improve service	Forecasts embedded in the reports people already trust get acted on in the same workflow, not filed as a separate deliverable.
Position for future	For many organizations this is the on-ramp to ML, and where versioning, monitoring and retraining discipline is first established.

PERFORMANCE DIMENSION	ILLUSTRATIVE TARGET	OUR CURRENT STATE
Forecast accuracy, per use case	Meets or beats an agreed baseline: naive trend projection or prior-year comparison	
Forecast monitoring	100% of production forecasts tracked against actuals on a defined cadence	
Retraining cadence	Defined and adhered to; drift past a threshold triggers retraining rather than waiting for review	
Integration into decisions	Share of forecasts consumed in the reporting layer, not sitting in a standalone output	

ASK BEFORE YOU SET A TARGET

- Are today's forecasts tracked against actuals at all? If this month missed by 30%, would anyone know before next quarter's planning cycle?
- Does the forecast live in the governed semantic layer, or as a one-off notebook model drifting away from everyone else's metric definitions?
- Who owns the model now that the analyst who built it has moved on?

CAPABILITY 4

Prescriptive Optimization

The ability to recommend or automatically execute the best available action given a forecast, constraints and an objective: what-if modeling, optimization engines, and the automated or human-in-the-loop mechanisms that turn a recommendation into an executed action, with the outcome fed back to confirm it worked.

The most organizationally demanding capability in the layer: skip trustworthy forecasts, explained context or a governed baseline and recommendations have nothing reliable to optimize against.

Protect revenue	Dynamic pricing and next-best-action recommendations shape transaction value in the moment, not after the fact.
Reduce cost	Automated reorder points and staffing optimization cut the cost of overstock and stockouts at the same time.
Improve service	What-if tools let decision-makers pressure-test a lever (price, budget, headcount) before committing to it.
Position for future	Closing the loop, measuring whether the recommended action improved the outcome, is what separates a one-off project from a capability that compounds.

PERFORMANCE DIMENSION	ILLUSTRATIVE TARGET	OUR CURRENT STATE
Recommendation-to-action rate	Share of recommendations actually executed, not generated and ignored	
Outcome validation	Share of executed recommendations with a measured before/after outcome fed back in	
What-if adoption	Share of relevant pricing, staffing and budget decisions run through scenario modeling first	
Guardrail coverage	100% of automated actions have a stopping condition and a human override path	

ASK BEFORE YOU SET A TARGET

- Of last quarter's recommendations, how many were executed, and of those, how many were confirmed to have worked? If neither is tracked, the gap is the feedback loop, not the optimization logic.
- Are the engine's constraints and objective still the business's current goal, or were they hardcoded for last year's conditions?
- Which decisions are safe to automate, and which stay human-reviewed regardless of technical feasibility? Without that decision, automation scope creeps ahead of readiness.

Four things that hold across all four

Trust compounds upward, and so does its absence.

A prescriptive recommendation built on an unmonitored forecast, built on an unexplained anomaly, built on a dashboard nobody trusts, inherits every one of those weaknesses at once. Fixing distrust at the prescriptive layer almost always means fixing a descriptive-layer gap first, not adding optimization sophistication.

Value shifts from informing to acting as capability matures.

Descriptive and diagnostic work mostly improves service and cost by making the truth cheaper and faster to find. Predictive and prescriptive work starts shaping revenue and risk directly, because its output changes what the business does rather than what it knows. That shift in stakes should come with a matching shift in sign-off rigor: a refresh failure and an automated pricing error are not the same class of incident.

Risk changes in kind, not just in degree.

Descriptive risk is mostly staleness and inconsistent definitions. By prescriptive optimization it is an automated action reaching a real customer or a real balance sheet with no human checkpoint. One review process applied uniformly across all four will slow low-stakes dashboard work down and still under-govern the highest-stakes actions.

A gap is only as good as the current-state number behind it.

Every illustrative target here needs a real current-state figure before it becomes actionable. The fastest route to usefulness is to fill in the current-state column for the one or two capabilities under active investment this planning cycle, not to complete all four at once.

DO THIS NEXT

Pick the single capability under active investment or debate right now. Complete its performance table with real current-state numbers from whatever monitoring exists today, even informal monitoring. That filled-in gap, not the illustrative targets above, is what should drive the next round of prioritization and investment conversations.