

Statistical Engine Tuning for improved forecast accuracy

Make a significant impact on the effectiveness and credibility of Oracle Demand Management and your commercial bottom line with a Kalypso Statistical Engine Tuning.

THE CHALLENGE

We know that poor forecasts lead to many inefficiencies including out-of-stocks, material shortages, higher inventories, and over-compensation via safety stock policies.

Statistical forecasts are a key component of most demand planning implementations, and usually serve as the default forecast before manual overrides are made. Engine Tuning is the process of deeply examining clients' demand patterns and adjusting the levers to improve their statistical forecast accuracy

Poor statistical forecasts lead to a variety of pathologies, including higher forecast errors, excessive time spent on manual overrides, planner-driven forecast bias, unproductiveness, user frustration and so on. But *well-tuned* statistical forecasts lead to lower forecast errors and leave more time for planners to engage in productive activities such as reviewing and investigating forecast *exceptions* for high-value items.

In both Demantra and Demand Management Cloud (DM Cloud), Oracle provides a sophisticated and very configurable “machine learning” statistical forecast engine. The impressiveness of the engine should drive us to want to exploit it to the fullest. It has a range of levers to influence the shape of forecasts and optimize them for each individual client's demand patterns.

These include:

- Ability to model **causal factors** that influence demand such as trend, seasonality, special holidays, events, end-of-quarter patterns and so on (and for Demantra, promotion uplifts and activity shapes). Some causal factors are pre-configured (but can typically be modified) and others can be designed and created for a particular client.
- **Forecast tree** configuration – refers to the levels of aggregation of history data that the engine considers when generating forecasts
- Selection from around 15 different **engine models** – various flavors of Regression (these use the causal factors), plus Crostons, Holt, etc
- Settings of many business-related statistical **parameters**, such as those that impact forecast validation, outlier detection, intermittency detection, naïve forecasts, etc

OUR SOLUTION

1. To help organizations understand the potential improvements, we provide an **Engine Tuning Assessment** service. In this assessment we agree on the details of error metrics calculations that suit your business and establish a “baseline” weighted forecast error and bias for the business. We show you which combinations are driving the worst weighted forecast error; we look at a variety of demand patterns / anomalies; we look for forecasts that are clearly misaligned with historical demand or misaligned with the business's view of “what is a good forecast”; and we discuss the drivers of demand and what kinds of remedies are possible through a full Engine Tuning exercise.
2. In the full Engine Tuning engagement, we perform the above “assessment” activities followed by iteratively working through all identified issues. We use “simulation” to check localized impacts of engine configuration changes, and (typically) perform full engine runs after each change to assess the impact on the whole data set – using the metrics established earlier. We publish and walk through a “tuning configuration” report at the conclusion. We work closely with the business through the entire process, including a period for post-go-live Support.

BUSINESS-ORIENTED

Kalypso's Engine Tuning is not a technical service that happens in a vacuum. It addresses configuration of all the levers mentioned above, given:

1 Your business's historical demand profiles:

- i.e., seasonality, trends (short and long term), intermittency, stability or lack of it, life-cycle length, holidays that impact demand, other causality (i.e. promotions, end-of-quarter)
- Historical disruptions like product availability shortages, COVID-19 impact

2 Taking into consideration your knowledge of your own markets:

- Is there cannibalism between new/old products or between different locations (i.e., ship-from)
- Will historical trends likely continue?

3 And your own forecasting preferences:

- Conservativeness: should forecast trend/level react slowly or quickly to recent changes in demand?
- At what base-level should forecasts usually be generated?
- When should combinations be considered "dead" (due to no recent demand), so that no further statistical forecast is generated?
- How aggressively should we try to detect outliers, if at all?

Each business is different, with different influences on demand, and varying degrees of randomness of the fluctuations in demand (which cannot be modeled/predicted). And there are nuances in the way forecast error is measured in each business. Our focus is on improving forecast errors during the tuning exercise.

WHEN IS THE RIGHT TIME TO TUNE?

Engine tuning is a standard part of Kalypso's implementation methodology for new Oracle demand management projects. We recommend tuning before Go-Live, if possible, between SIT and UAT. But any time is a good time to consider investing in your first tuning exercise. A well-tuned statistical forecast creates a good impression, increases productivity, drives down forecast error, and fuels user acceptance.

Markets change over time, and different influences on demand become apparent. For this reason, and especially if statistical forecast accuracy begins to wane, you should consider a refresher tuning exercise every two years. This helps to keep the demand planning system fresh, and the consultative exercise will help any new demand planners become more productive and confident with their Oracle system (and less likely to regress back to Excel.)



Andrew Calder
Technical Manager, Kalypso
Andrew.calder@kalypso.com

WHY KALYPSO?

Kalypso has a careful, consultative approach to engine tuning. We have conducted around 40 tuning exercises in the last seven years, with every single engagement showing empirical improvements. Our engine tuning engagements typically drive weighted forecast error improvements between 10-20% (relative to the existing error percent), with some much better than this. Vast improvements in overall forecast bias are very common.

Deep knowledge of the Oracle forecasting engine, understanding causality, avoiding over-fitting, curiousness, attention to detail, experience with different demand profiles – these are all necessary requirements to conduct an Engine Tuning. Kalypso brings all these. But we don't treat tuning as just a technical exercise. We value input from the businesses we engage with and work hard to transfer knowledge to users to make the Oracle engine a little less of a black box. Clients often gain new insights into their own data as they are forensically analysed to help produce the best forecast.

Make a significant impact on the effectiveness and credibility of your Oracle Demand Management application, and your commercial bottom line with a Kalypso Statistical Engine Tuning.