



10 Tips for Implementing SAP APO DP

- Top Five System Design Tips
- Top Five Forecasting Tips



Spinnaker is a client advocate and is in no way sponsored by any software vendor.

The success of your APO DP installation will be defined by the people who use the system. If users are satisfied, you'll have achieved success. And success will depend heavily on your early efforts to develop a solid, well thought out design.

The SAP Advanced Planner and Optimizer software suite is the centerpiece of SAP's Supply Chain Management solution.

In the several years since its introduction, APO has enjoyed increasing popularity among users, with each subsequent release marked by functionality enhancements and improvements in existing features.

There are five distinct modules that comprise the APO suite:

- Demand Planning (DP)
- Supply Network Planning (SNP)
- Production Planning and Detailed Scheduling (PP/DS)
- Global Available to Promise (GATP)
- Transportation Planning/Vehicle Scheduling (TP/VS)

Of these five modules, Demand Planning is the one most frequently installed by supply chain managers, thus it's the module that most frequently poses challenges for implementers.

Based on our experience over the years, installing and configuring DP for a number of leading multinational corporations, we've identified **10 core issues likely to have the greatest impact** on the success of DP implementations. And so we compiled the following roster of Top 10 Tips for addressing each of these key implementation issues, along with corresponding recommendations for driving the best possible

outcomes, to provide helpful insight for anyone involved in a first-time APO DP implementation.

Before you dig into the tips, it's worth noting that all 10 fall within one of two categories: those issues related to non-forecast functionality — including sound, fundamental design issues based on widely accepted best practices; and those issues related purely to APO forecasting functionality.

Tips One through Five— System Design

1. Try to limit your design to a single planning area.

Planning areas constitute the central data structures of APO DP (and SNP), so before you can begin planning or forecasting in APO DP ***you must first define a planning area.***

A planning area specifies the unit of measure on which all forecasting will be based, as well as corresponding parameters like currency, the size of time buckets in which data will be managed (days, weeks, months, etc.), and key figures like Orders and Forecasts, in which data will be stored (see **Figure 1**).

Of course, planning areas require maintenance, and having multiple planning areas will compound your maintenance responsibilities, but in some cases, you may be forced to define multiple planning areas.

Any trademarks referenced here are the property of their respective owners.

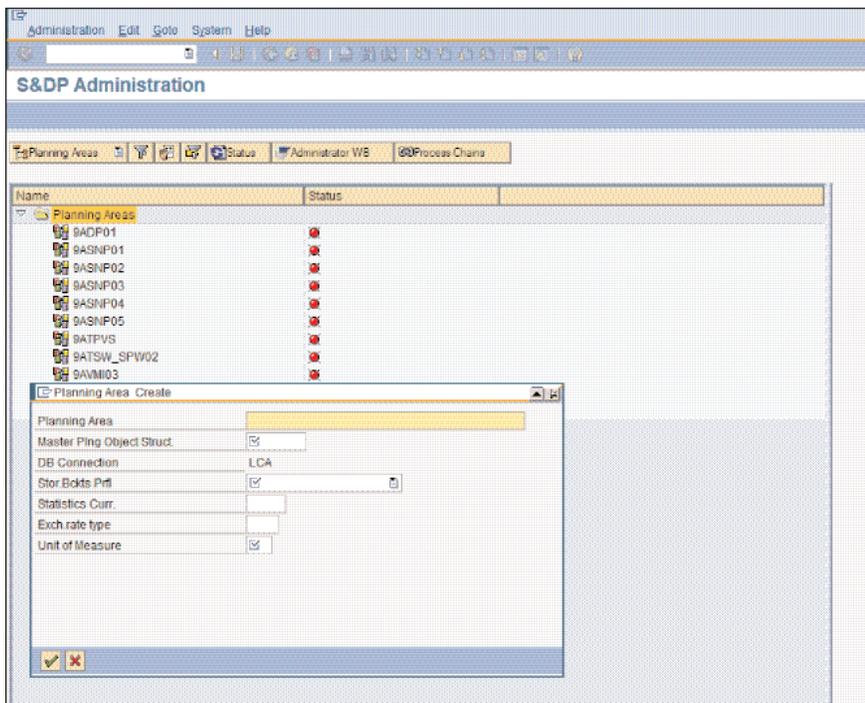


Figure 1 — Creating a Planning Area in APO DP

For example, if your business requires planning of products that don't share a common unit of measure — or units that can't readily be converted from one measure to the other, like pounds and square feet — you'll need to set up separate planning areas for each type of product. Other situations aren't as clear cut.

For instance, you would also need to define multiple planning areas in a business scenario where planning for the **majority** of products is done over a one-year horizon, on a weekly basis, and at a detailed level, like product-customer, but where **longer-term planning** is done over a multi-year horizon, in monthly buckets, and at an aggregated product level, like product family or brand.

Since data in APO DP is always stored at the lowest level of granularity, which in this latter case would involve weekly buckets, it may be better to design separate planning areas so that people working at an aggregate level wouldn't be burdened by slower system response times, necessitated by those people working at a more detailed level.

In such a case, it would be worthwhile to create one planning area for regular use by the demand planners and a separate, more streamlined planning area for the aggregate planners, to ensure optimal system performance for both groups.

The point is, from the very outset of the design phase, it pays to carefully evaluate your need for multiple planning areas.

2. Use navigation attributes.

In terms of ensuring a successful Demand Planning implementation design, the master planning object structure (POS) is even more fundamentally important than planning areas.

The POS contains the characteristics that enable planners to slice and dice data into groupings that are meaningful from a business perspective, such as geographical areas (Europe, northeast); business locations like plant identifiers; business segments (e.g., sales organization, division, etc.);

The benefits of using navigation attributes outweigh the drawbacks; you'll find they pay long-term dividends in terms of improved system performance.

customer designations; and product hierarchy levels.

And when creating your master planning object structure, consider using navigation attributes as an alternative to characteristics. One of the biggest advantages of this approach is that navigation attributes help you minimize the number of characteristic value combinations.

CVCs represent objects that are planned in Demand Planning. These objects are formed by various permutations of the characteristics that system managers may specify in the POS. For example, one CVC may represent a particular product shipped from a particular plant to a particular customer in a particular location.

Clearly, the more characteristics there are, the more CVCs you're likely to end up with. And while it's not uncommon to have hundreds of thousands of CVCs in a Demand Planning installation, **the more CVCs, the greater the drain on system performance.**

Upgrading to a computer platform with more memory is one way to overcome this bloat, but defining navigation attributes at the outset of system design helps you prevent bloat from occurring in the first place.

Unlike characteristics, which represent a specific quality, such as a location name or a product code,

navigation attributes derive their values dynamically from master data that resides in the system. And since master data is subject to change, using "virtual" navigation attributes means that no CVC is actually created for storage, thus reducing the overall system processing burden and consequently improving system performance.

Of course, there can be drawbacks to using navigation attributes, such as their limitations for use in fixed aggregates and promotion planning (the trade-offs are discussed in the SAP online knowledge base, OSS Note 413526). However, the drawbacks are limited to specific aspects of certain functions in DP, and it's been our experience that the benefits of using navigational attributes outweigh the potential negatives. If you don't expect your use of DP to be limited by these potential drawbacks, or if you decide the trade-offs are worthwhile, **you'll find that using navigation attributes will pay long-term dividends in terms of improved system performance.**

3. Be careful of fixed aggregates.

Fixed aggregates offer a convenient means of grouping data elements that are typically or frequently viewed collectively.

As shown in **Figure 2**, APO stores data at the lowest possible level of detail, but when you define a fixed aggregate, like **Plant 1** or **REGION1**,

The three most important factors for ensuring the integrity of an APO system are:

- master data,
- master data,
- master data.

data for that collective element is also stored in an aggregated form.

This helps speed the display of data in interactive planning, but the operative word here is *display* of data, since using fixed aggregates can significantly slow down system response time when *saving* data, in the course of everyday planning.

Before the system can display a fixed aggregate, it must first assemble all the associated data components. Similarly, because of the redundant nature of fixed aggregates, simply making a change in data could trigger the system to process two or more changes, to populate the new data throughout every associated fixed aggregate. This is the same reason it takes longer to create time series objects when fixed aggregates are involved.

Since it's not advisable to modify fixed aggregates *after* activating your system's planning object structure, it's best to evaluate your need for them while still in the design phase of implementation. (Refer to OSS Note 503363 or more information.)

4. Don't underestimate the difficulties of migrating data from a legacy system.

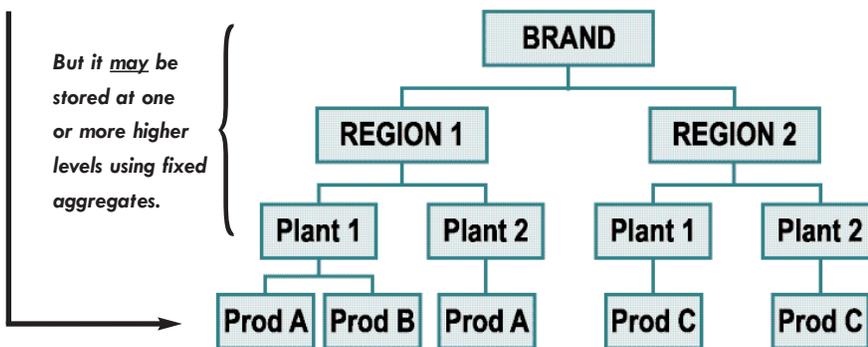
To borrow from the old adage about the three most important factors to consider before buying real estate — location, location, location — it's worth noting that the three most important factors for ensuring the integrity of an APO system are master data, master data, and master data.

This principle applies especially to the initial setup and activation of an APO Demand Planning environment, but it's also crucial to the success of ongoing demand planning operations.

Master data can mean different things in different systems, even within APO. In SNP, master data is more like R/3 master data, but in APO DP, master data normally refers to the CVCs generated in the master POS. In terms of system design and implementation, we're referring more to the latter definition, but not ignoring the former. The point is, regardless of what type of data must be moved from one system to another, unforeseen problems are the rule rather than the exception.

To ensure success, assign your best people to this task, make sure you have strong BW support (Business Information Warehouse), and allow *adequate* time — more than you think you'll need — to accomplish the task.

Figure 2 — Data is stored at the lowest level by default.



***What you SEE
may not be what you GET.***

Take care when using Event macros — the data you see onscreen may not match the data stored in liveCache.

5. Proceed with caution when using Event macros.

Macros provide a handy, convenient means of automating calculations that you perform repeatedly during planning activities; and APO has its own macro language you can use to create, save, and store mathematical shortcuts, to manipulate planning data at your convenience or automatically.

The term “Event” macros refers collectively to macros that execute automatically in a planning book data view within interactive planning.

As the name implies, Event macros are triggered by a specific *event*, such as when you drill upward or downward through data, or when you load data into the planning table — or even simply press the <Enter> key — during an interactive planning session.

Event macros are a terrific convenience for manipulating data, but if you’re not careful how you use them, they can sometimes contribute to a misrepresentation of facts.

For example, when you use an event macro to alter data in your data view, the modified values you see onscreen are no longer, in fact, the same as those stored in system memory, which is known in APO as liveCache. Onscreen values only become actual stored values in the system after you save them.

System prompts may serve as reminders to help avoid confusion between data stored in liveCache and data displayed onscreen. You can also create homegrown warnings and alerts in the data view (perhaps triggered by other event macros) to help avoid inconsistencies.

Regardless of what you do, this issue can be a tricky one, so take care to recall that what you *see* onscreen *may not be* what you *get* — the *real* data stored in memory.

Tips Six through 10 — Forecasting

6. One size *does not* fit all.

APO DP provides some wonderful tools for managing data, but these tools can be misused all too easily. Here’s an example:

One of the forecasting strategies offered in DP — strategy 56 — is called Automatic Model Selection II. It automatically selects an appropriate forecasting model based on your historical demand curve. But don’t expect a single forecast model to be right for all products.

Many first-time users of APO view strategy 56 as an answer to their prayers — a forecasting autopilot. But just as the autopilot system on an airplane can’t be used at all times in all situations, from takeoff to landing, so too must different solutions be applied to different circumstances when forecasting.

Various demand patterns, items at different stages of the product life cycle, high volume vs. low volume — these are just a few of the reasons why one size does not fit all when it comes to forecast models — even one that purports to automatically select the right strategy.

7. Don't try to justify — or dismiss — forecast results by “eyeballing” historical data.

The most frequent complaint we hear from planners and supply chain managers about APO forecast results is that they're not “reasonable.”

Many users, regardless of job title, industry vertical, or skill level in the planning realm, *insist* their data exhibits specific traits, like seasonality, for example, so they can't understand why their DP model would produce a constant forecast (see **Figure 3**). Some users will even point out what they perceive as

seasonal patterns in their historical data. That's because people are often more inclined to believe their eyes than their data. In most cases like this, though, the “patterns” that people see either aren't statistically significant, or there's some other characteristic of demand history, not readily apparent to the human eye, that negates the initial impression.

Sometimes, however, users *are* right, and models *don't* return the best results; this could be caused by any number of factors, primarily:

- One or more incorrect settings or parameters in the model
- Optimization of an error measurement that may not be appropriate for the business
- Choice of the “wrong” strategy

The key here is to overcome the human tendency to be defensive and instead focus on solving the problem.

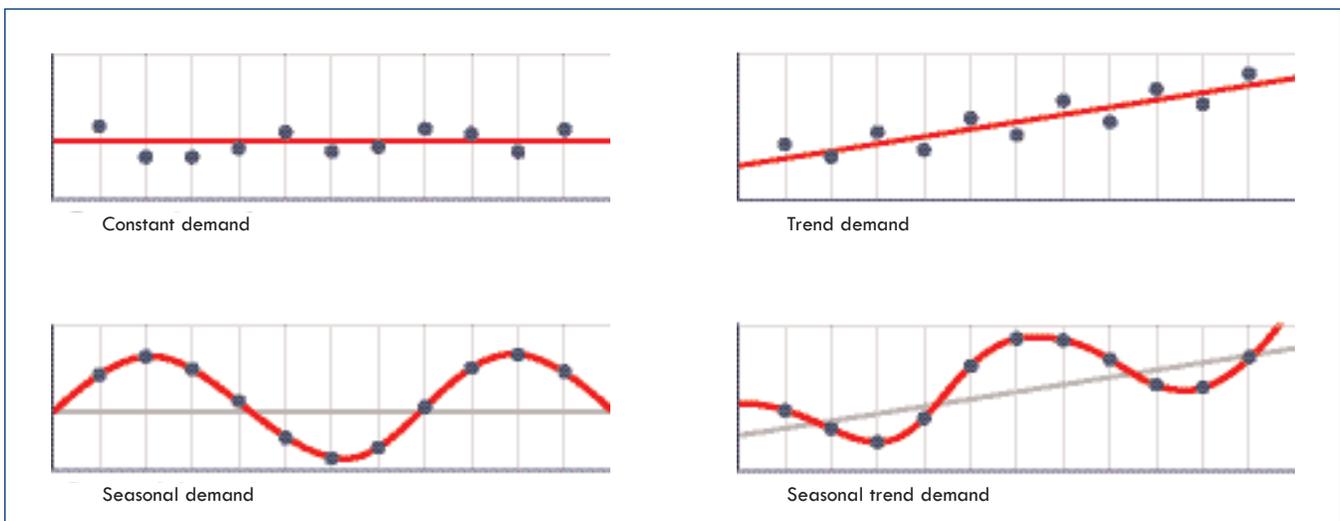


Figure 3 — Typical demand patterns: Constant, Trend, Seasonal, and Seasonal Trend

No planning system will return flawless results for your business, but if you are **truly confident** that your data exhibits seasonality (to continue using the previous example), the best approach is to try another forecast strategy — a seasonal model, for example — to “force” seasonality on your data rather than rely on an automatic model selection strategy that may or may not recognize seasonality.

Take time to consider your specific objectives up front. And be mindful of these 10 tips.

They're time-tested, field-proven ways to get the most value out of your APO Demand Planning installation.

8. Don't try to compare APO forecast results to those from other software.

Another frequent comment we hear — sometimes as a complaint, sometimes as a puzzled observation — is that APO “didn't return the same forecast result as our old software,” even though the type of forecasting approach used was the same, and the historical data on which the forecast was based was the same.

Obviously, if both software applications used and applied identical algorithms in exactly the same way, then the forecast results would be the same. But this rarely happens in practice. In addition to slight differences in mathematical approaches to certain statistical

problems, there are numerous parameters that can vary, both in value and in the way they're applied. Some parameters may be varied internally by software, as a means of optimizing calculations, for example; others may be set by the user.

In any case, different forecasting packages use different approaches, so similar results are likely to be the exception rather than the rule.

9. Use higher gamma factors.

If your products exhibit seasonality, then you'll need to pay close attention to the gamma factor, the parameter within statistical forecasting that gives planners the ability to smooth seasonality exhibited by data.

Note that the gamma factor does not *determine* whether or not a product demonstrates seasonality; it determines how pronounced you want the seasonality to be reflected in your forecast (the same way the beta factor impacts trend).

Accordingly, the gamma factor is not determined by the data; it is set by a forecaster. The default gamma factor within APO DP is 0.3. This is a reasonable value and in line with generally accepted practices, however, we've found that setting higher gamma factor values can result in more accurate forecasts for many companies (see **Figure 4**).

Figure 4 —
When creating and defining a univariate forecast profile within APO DP, don't be afraid to try higher gamma values.

The screenshot shows the 'Univariate Forecast Profile' configuration window. The 'Model Parameters' section is expanded, showing the following settings:

- Forecast Strategy: 56 Forecast with Automatic Model Selec
- Alpha: 0.20
- Beta: 0.30
- Gamma: 0.50 (highlighted with a red box)
- Sigma: (empty)
- Periods: 12
- Variation: (empty)
- Alpha 2: (empty)
- Error Measure: MAPE
- Alpha Start Value: (empty)
- Alpha End Val.: (empty)
- Alpha Increment: (empty)
- Beta Start Value: (empty)
- Beta End Value: (empty)
- Beta Increment: (empty)
- Gamma Start Value: (empty)
- Gamma End Val.: (empty)
- Gamma Increment: (empty)
- Weighting Profile: (empty)
- Trend Dampng Prfl.: (empty)
- Hist. Val. Markngs: (empty)
- Diagnosis Group: (empty)

The 'Control Parameters' section includes:

- Outlier Correction: Ex-Post Method
- Without Leading Zeros:
- Add Up Decimal Places:
- Days in Period: (empty)

The 'Forecast Errors' section includes:

- MAD:
- MSE:
- RMSE:
- MAPE:
- MPE:
- Error Total:

The 'Promotion' section includes:

- Key Figure: (empty)
- Select:
- Change Vals:

If your historical data exhibits seasonality, don't be afraid to experiment with higher gamma factors. You may be pleasantly surprised by the results.

10. Use composite models.

Admittedly, you could build a career around APO by experimenting with individual forecast profiles and all the settings and parameters that go with them. So why are we suggesting that planners use composite models? Because they're a powerful tool, too often underutilized by APO practitioners.

Composite models are defined simply as the combination of two or more individual forecast profiles. You create them by combining one or more univariate profiles (also known as time series profiles) and/or one or more multiple linear regression models (also known as causal models). Each profile is assigned a weighting factor, which is simply the percentage of the profile's contribution to the final forecast result.

It's true that the number of these possible combinations could literally be endless, but view this as a positive, not a negative. You don't *have* to try every combination (indeed, you can't), and constructing a composite model that improves your results — often significantly — is usually not as overwhelming as you might think.

If you've never experimented with composite forecast models in APO, you'll find the functionality quite user friendly (see **Figure 5**).

Parting Thoughts

The ten recommendations we've detailed here are time-tested, field-proven ways to get the most value out of your APO Demand Planning installation efforts.

We **strongly recommend** that you consider Tips 1-5 from the very outset of your blueprinting process and during the planning stages of your implementation.

Focus on the remaining tips as part of the ongoing learning curve that will evolve following the go-live phase.

And finally, consider these two bonus recommendations:

11. Get help.

Securing expert guidance at the outset of your implementation can head off the need for regrouping down the road. And the farther down the road you go in terms of implementation, the costlier it will be to get back on track if you veer off course.

12. Design your solution with the end goal in mind.

Each step of the way, think through the details of what you’re trying to accomplish. You’re more likely to achieve success when you identify specific objectives at the start of a project.

In the end, the success of your Demand Planning installation will be defined by those who use the system. If users are satisfied with its performance and happy with its results, then you’ll have achieved success. And that level of success will depend heavily on your up-front efforts to develop a solid, well thought out design concept, based firmly on many of the recommendations presented here.

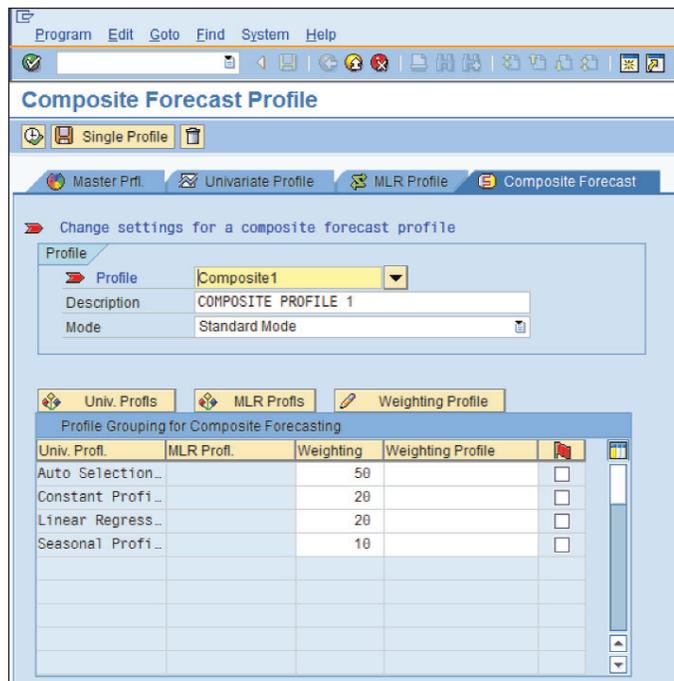


Figure 5 — Defining a composite forecast profile

#



To learn more, contact the supply chain specialists at Spinnaker. Call **877-476-0576** or visit **www.spinnakermgmt.com**.