



Microsoft Dynamics AX 2012 R3

Demand forecasting

White paper

Roxana Cristina Diaconu and
Mehmet Yurdal
May 2014

This white paper explains the value proposition of the Demand forecasting tool in Microsoft Dynamics AX 2012 R3 and gives an overview of the functionality.

www.microsoft.com/dynamics/ax

Send feedback.

Contents

Demand forecasting	4
Value proposition	4
Basic flow	5
Demand forecasting setup	5
SQL Server Analysis Services parameters setup	10
Generating a statistical baseline forecast	11
Making manual adjustments to the baseline forecast	13
Handling the Excel issue with weighted allocation when cells that have a value of 0 are modified	16
Making percentage adjustments to the forecast	19
Scenario 1: Every cycle, make an optimistic demand forecast that overall demand will increase by 10 percent across all companies and products	19
Scenario 2: Every cycle, make an optimistic demand forecast that overall demand will increase at a specified rate each month across all companies and products	21
Scenario 3 – Phase-out: For a product that is reaching the end of its life, decrease the forecast by a certain percentage each month	24
Scenario 4 – Phase-in: Create a forecast for a new product by copying the forecast from another product and adjusting it by a certain percentage	26
Importing an adjusted forecast	28
Monitoring forecast accuracy	29
Monitoring forecast generation and import	30
Understanding the Analysis Services forecasting methods	31
AX 2012 R3 DemoData	31
Demand forecasting approaches	32
Introduction to Analysis Services time series forecasting	32
The ARTxp forecasting algorithm	34
The ARIMA forecasting algorithm	34
Demand forecasting parameters	34
AUTO_DETECT_PERIODICITY	34
PERIODICITY_HINT	34
FORECAST_METHOD	37
INSTABILITY_SENSITIVITY	39
PREDICTION_SMOOTHING	40
MINIMUM_SERIES_VALUE	43
MAXIMUM_SERIES_VALUE	43
MISSING_VALUE_SUBSTITUTION	44
COMPLEXITY_PENALTY and MINIMUM_SUPPORT	48
Tips for setting up demand forecasting algorithm parameters	48
References	49

Performance improvement tips and tricks	49
A brief technical look at the Demand forecasting implementation	49
Hardware recommendations	49
System setup	50
SQL Server and Analysis Services editions	50
A deeper look at the steps for statistical forecast generation	50
Performance improvement tips for forecast generation	51
Demand forecast performance cheat sheet	52

Demand forecasting

Microsoft Dynamics AX 2012 R3 Demand forecasting is a lightweight yet powerful tool that lets users predict demand based on historical data, adjust the predicted demand, and import the values into Microsoft Dynamics AX forecast models.

Demand forecasting provides high return on investment (ROI) in a short time by taking advantage of familiar Microsoft technology. To generate the baseline forecast, we rely on the power of Microsoft SQL Server Analysis Services. To visualize the forecast, adjust the forecast, and view or create new key performance indicators (KPIs), we rely on the extensive functionality of Microsoft Excel.

Demand forecasting can be used to predict both independent demand (from sales orders) and dependent demand (at any customer order decoupling point). In combination with the enhanced forecast reduction rules in AX 2012 R3, this becomes an ideal solution for mass-customization.

Key features of Demand forecasting are:

- Generation of a statistical baseline forecast based on historical data. The forecast is expressed in quantity only.
- A dynamic set of forecast dimensions.
- Visualization of demand trends and adjustment of the demand forecast.
- Outlier removal.
- Forecast accuracy KPIs.

Value proposition

Because Demand forecasting is fully integrated with AX 2012 R3, there is no need for data migration and integration, provided that you keep your historical transactions in the AX 2012 R3 database.

Three major themes were followed during the implementation of Demand forecasting:

- **Modularity** – Demand forecasting is a modular and easy-to-configure tool. You can enable and disable the functionality by toggling the configuration key under **Trade > Inventory forecast > Demand forecasting**.
- **Reuse of the Microsoft stack** – Microsoft Dynamics AX partners and independent software vendors (ISVs) already have resources that can work on the Microsoft stack – Microsoft SQL Server 2008 and later versions, Microsoft SQL Server Analysis Services 2008 and later versions, and Microsoft Excel 2009 and later versions. This makes Demand forecasting easy to extend and adapt to user needs.
Analysis Services offers modern and highly accurate time series algorithms for demand forecasting. For more information about the forecast methods and the parameters they take, see [Microsoft Time Series Algorithm Technical Reference](#) and also the [Understanding the SQL Server Analysis Services forecasting methods](#) section of this white paper.
- **High ROI in a short time** – Demand forecasting is available for companies that have acquired a Production Planner user license. It does not require any additional software components besides the standard Microsoft Dynamics AX prerequisites.
Users who are running AX 2012 R3 on Microsoft SQL Server Standard Edition 2008/2012 will not have to change to a higher edition of SQL Server to use Demand forecasting. However, there are certain performance benefits if companies run SQL Server Enterprise Edition or SQL Server BI Edition. These come from the fact that, in SQL Server Enterprise Edition or SQL Server BI Edition, the historical demand data can be split among multiple forecasting mining models, which are then processed in parallel.

AX 2012 R3 offers the capability to do forecast reduction at any decoupling point. Demand forecasting for Microsoft Dynamics AX 2012 builds upon this functionality, offering an opportunity to forecast demand based at any decoupling point. In fact, it provides the capability to forecast both dependent and independent demand.

It's important to understand not only how Demand forecasting can help Microsoft Dynamics AX users but also what its limitations are. Demand forecasting for AX 2012 R3 is not best-of-breed forecasting software. It is a tool that helps users conduct their forecasting processes. It offers the core functionality of a demand forecasting solution and is designed in such a way that it can be easily extended. It was developed in collaboration with many manufacturing Microsoft Dynamics AX clients and members of the Microsoft Dynamics Customer Council, but it might not be the best fit for

customers who operate in industries such as retail, wholesale, warehousing, transportation, or professional services.

Basic flow



Demand forecast generation starts inside AX 2012 R3. Historical transactional data from the Microsoft Dynamics AX transactional database is gathered and populates a staging table, which is later fed into an Analysis Services mining model. With minimum customization, users can plug various data sources in to the staging table, such as Excel files, comma-separated value (CSV) files, and data from Microsoft Dynamics AX 2009, provided that the master data (item names, units of measure, and so on) is the same across the various data sources. This lets users generate a demand forecast that takes into account historical data spread around multiple systems.

Inside the mining model, Analysis Services time series forecasting methods are used to calculate a baseline forecast. The parameters for these forecasting methods are managed in AX 2012 R3. The forecast and the historical data, and the changes made to the demand forecast in previous iterations, are then made available in a SQL Server data cube.

Excel is used to connect to the demand forecasting cube so that users can visualize and modify the baseline forecast. The manual adjustments that are made have to be published back to the cube before they can be imported into AX 2012 R3. Note that, because there is only one demand forecasting cube at any time, there is only one version of the forecast.

Finally, from AX 2012, the adjusted forecast can be imported back to the AX 2012 R3 transactional database in the forecast models.

Despite the many system components that need to interact, users always work in familiar environments: AX 2012 R3 (for forecast generation and import) and Excel (for forecast visualization and adjustment).

Demand forecasting setup

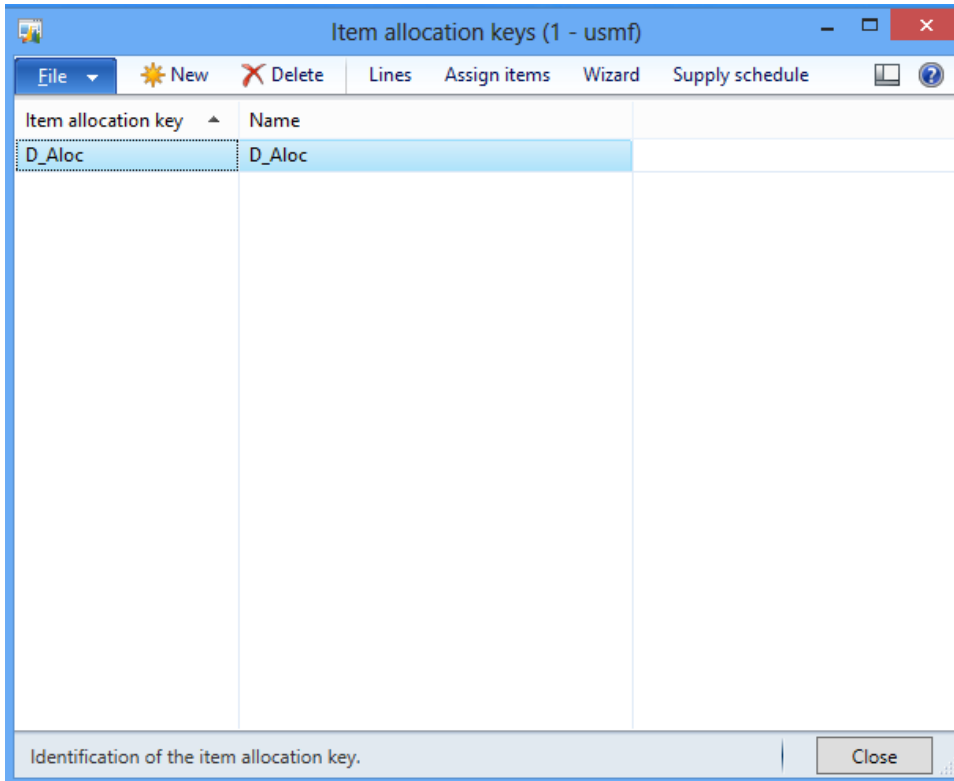
A demand forecast is calculated for a stock keeping unit (SKU; a released product and its dimensions) only if the SKU is part of an item allocation key. This rule was enforced to keep large numbers of SKUs manageable, and to enable SKUs that behave the same from a demand forecasting perspective to be grouped for performance reasons.

The item allocation key percentage is ignored at demand forecast generation time. The forecast is created based on historical data only.

An SKU must be part of only one item allocation key if that item allocation key is used in forecasting.

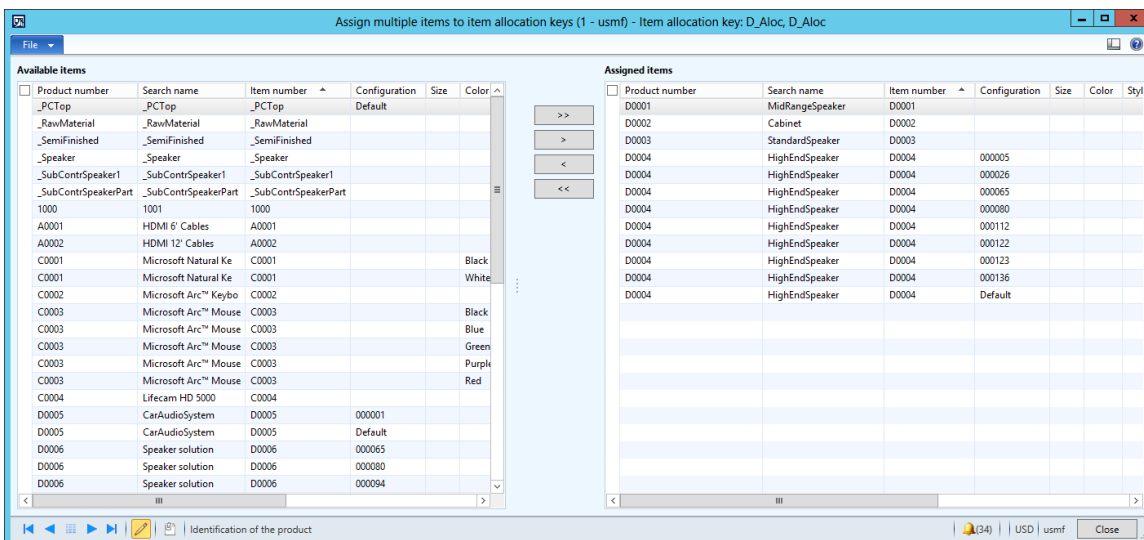
To add an SKU to an item allocation key, go to **Inventory and warehouse management > Setup > Forecast > Item allocation keys**.

The following illustration shows the **Item allocation keys** form.



As part of the Demand forecasting feature, we built a more user-friendly way to assign items to item allocation keys. Select the item allocation key that new SKUs should be added to, and then click **Assign items**.

The following illustration shows the **Assign multiple items to item allocation keys** form.



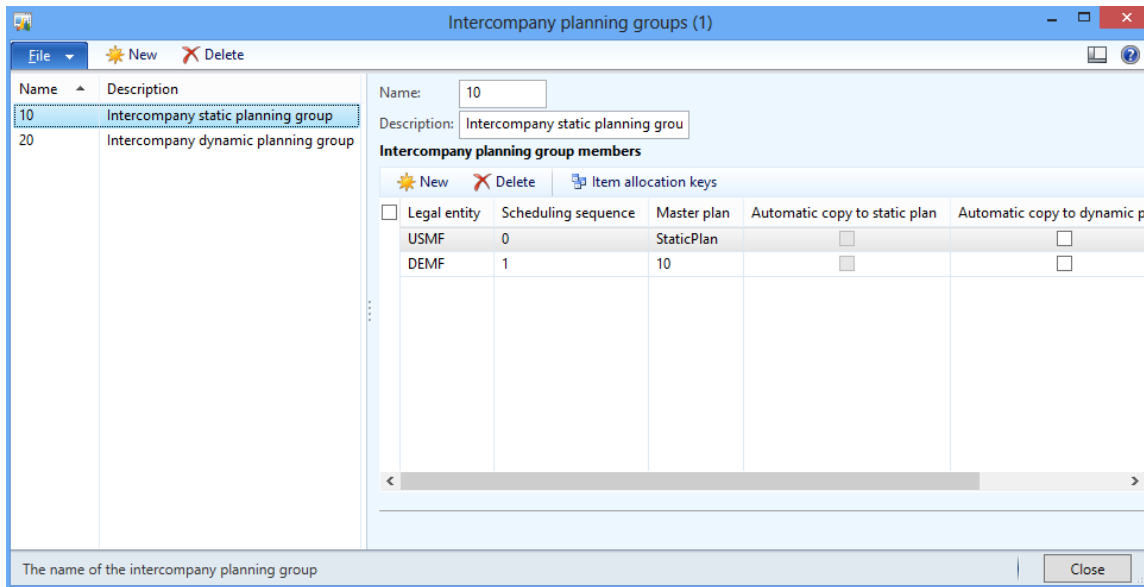
On the left side of the form is a list of all SKUs that are not assigned to the selected item allocation key. On the right side is a list of SKUs that are currently assigned to the item allocation key.

Demand forecasting generates cross-company forecasts. In AX 2012 R3, companies that are planned together are grouped into one intercompany planning group. To specify, per company, which item allocation keys should be considered for demand forecasting, associate an item allocation key with the intercompany planning group member by going to **Master planning > Setup > Intercompany planning groups**.

Note: If users don't want to generate cross-company forecasts, they can set a company filter at forecast generation

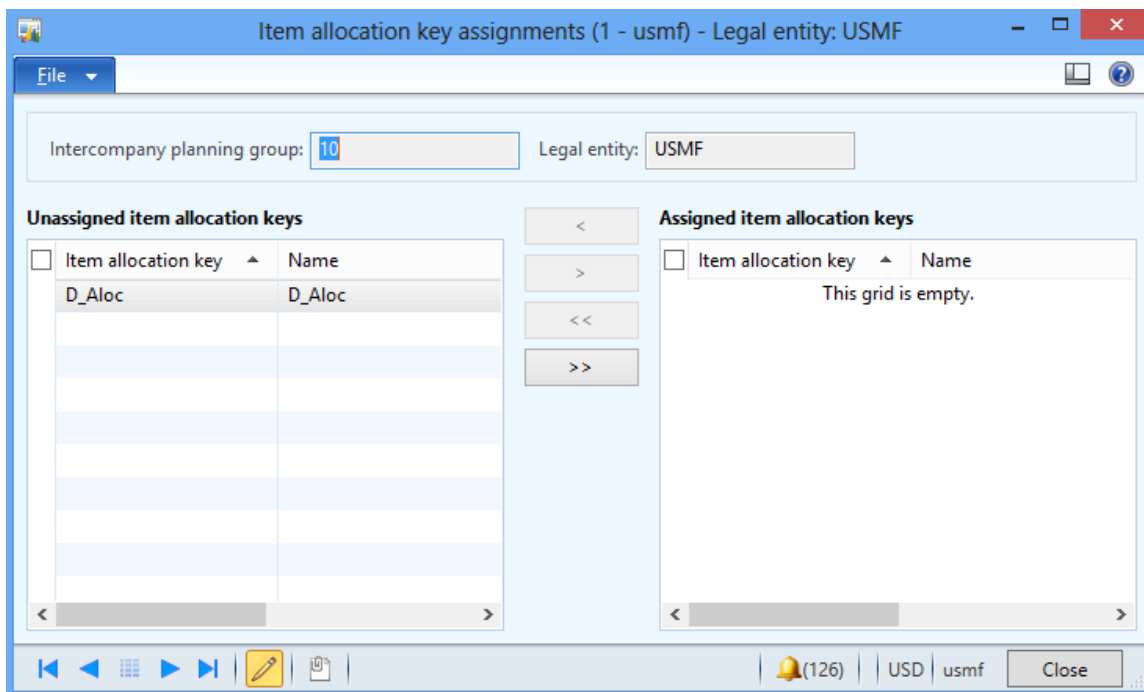
time. Alternatively, they can create an intercompany planning group with only one company, assign the item allocation key they want to generate forecast for to that company member, and filter by the intercompany planning group at forecast generation time.

The following illustration shows the **Intercompany planning groups** form.



Click **Item allocation keys** to add, remove, or modify the list of item allocation keys assigned to the intercompany planning group member.

The following illustration shows the **Item allocation key assignments** form.



The upper part of the form identifies the currently selected intercompany planning group (in this example, planning group 10) and the current intercompany group member (in this example, the USMF company). On the right side of the form is a list of all assigned item allocation keys from USMF. On the left side of the form is a list of unassigned item allocation keys from USMF. To assign item allocation keys to intercompany planning group members, or to unassign them, use the arrow buttons in the center of the form. To globally assign or unassign item allocation keys, use the

double arrow buttons in the center of the form.

By default, if no item allocation keys are assigned to intercompany planning group members, a demand forecast is calculated for all the SKUs assigned to all item allocation keys, from all AX 2012 R3 companies. Further filtering options for companies and item allocation keys are available in the **Generate statistical baseline forecast** form.

To set up Demand forecasting parameters, go to **Master planning > Setup > Demand forecasting parameters**. Because demand forecasting runs cross-company, the setup is global (in other words, it applies to all companies).

The following illustration shows the **Demand forecasting parameters** form.

The screenshot shows the 'Demand forecasting parameters (1)' window. The left sidebar has 'General' selected, with sub-items 'Forecast dimensions' and 'Item allocation keys'. The main area is titled 'Set up demand forecasting options'. It contains several input fields: 'Demand forecast unit' (ea), 'Demand forecast cube' (Demand Forecast initial), 'Demand forecast accuracy cube' (Demand Forecast Accuracy initial), and 'Excel demand forecast file path' (e:\temp\Forecast.xlsx). Below these is a section 'Include these historical transaction types to generate the statistical baseline forecast' with checkboxes for Sales order, Production line, Kanban job consumption, Kanban job transfer issue, Quotation, Project, Transfer, and Transaction. The 'Forecast algorithm parameters' section includes 'Forecast generation strategy' (SSAS time series forecasting methods) and a table for 'Forecast method used' with columns for Name, Value, and Explanation. The table has one row with 'Forecast method used', 'MIXED', and 'Options: ARTXP, ARIMA, MIXED'. At the bottom, there is a 'Close' button and a status bar that says 'The value of the parameter'.

Demand forecasting generates the forecast in quantities. Therefore, the unit of measure that the quantity should be expressed in must be specified in the **Demand forecast unit** field. The unit of measure must be unique to guarantee that the aggregation and percentage distribution in Excel makes sense. For more information about modifying aggregated values in Excel pivot tables, see the [Make manual adjustments to the statistical baseline forecast](#) section. For every unit of measure that is used for SKUs included in demand forecasting, make sure that there is a conversion rule that unit of measure and the general forecasting unit of measure. When forecast generation is run, the list of all items without unit of measure conversion is logged, so that users can easily correct the setup.

Additionally, the demand forecasting and demand forecasting accuracy cubes must be specified in the **Demand forecast cube** and **Demand forecast accuracy cube** fields. If the demand forecasting cubes have already been deployed, you can select the name in the lookup.

The path of the Excel file used to view and adjust the baseline forecast is specified in the **Excel demand forecast file path** field. Note that the value must be a file path, not a folder path. This can also be a Microsoft SharePoint path.

Demand forecasting can be used to forecast both dependent and independent demand. For example, if only the **Sales order** check box is selected, and if all the SKUs considered for demand forecasting are items that are sold, the system

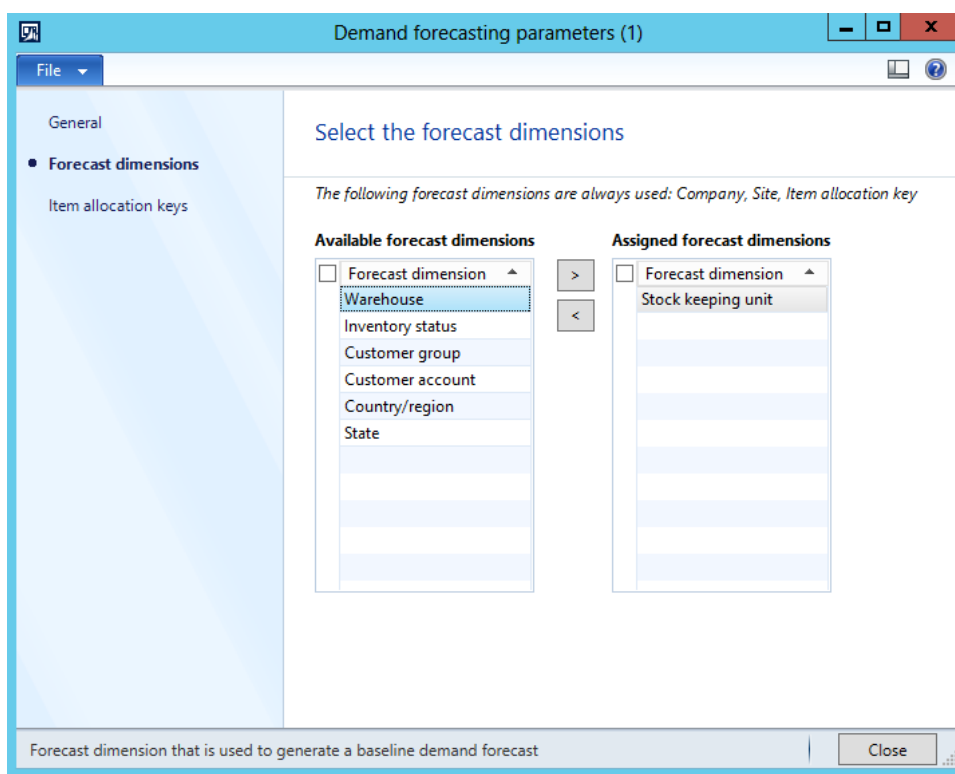
calculates independent demand. However, critical subcomponents can be added to item allocation keys and therefore considered for demand forecasting. Then, if the **Production line** check box is selected, dependent forecast is calculated.

There are two ways to create a baseline forecast in AX 2012 R3: by using forecasting models on top of historical data or by simply copying over the historical data to the forecast. The **Forecast generation strategy** field lets you select between these two methods.

The **Do not allow gaps in historical data** check box controls the way AX 2012 R3 handles gaps in historical data. For more information about this parameter, see the [Understanding the SQL Server Analysis Services forecasting methods](#) section.

By clicking the **Forecast dimensions** link in the left pane of the **Demand forecasting parameters** form, you can also select the set of forecast dimensions to use when the demand forecast is generated. A forecast dimension designates the level of detail for which the forecast is defined. Company, Site, and Item allocation key are mandatory forecast dimensions, but you can also select to generate forecasts at the warehouse, inventory status, customer group, customer account, country/region, state, and SKU levels.

The following illustration shows the **Forecast dimensions** area of the **Demand forecasting parameters** form.



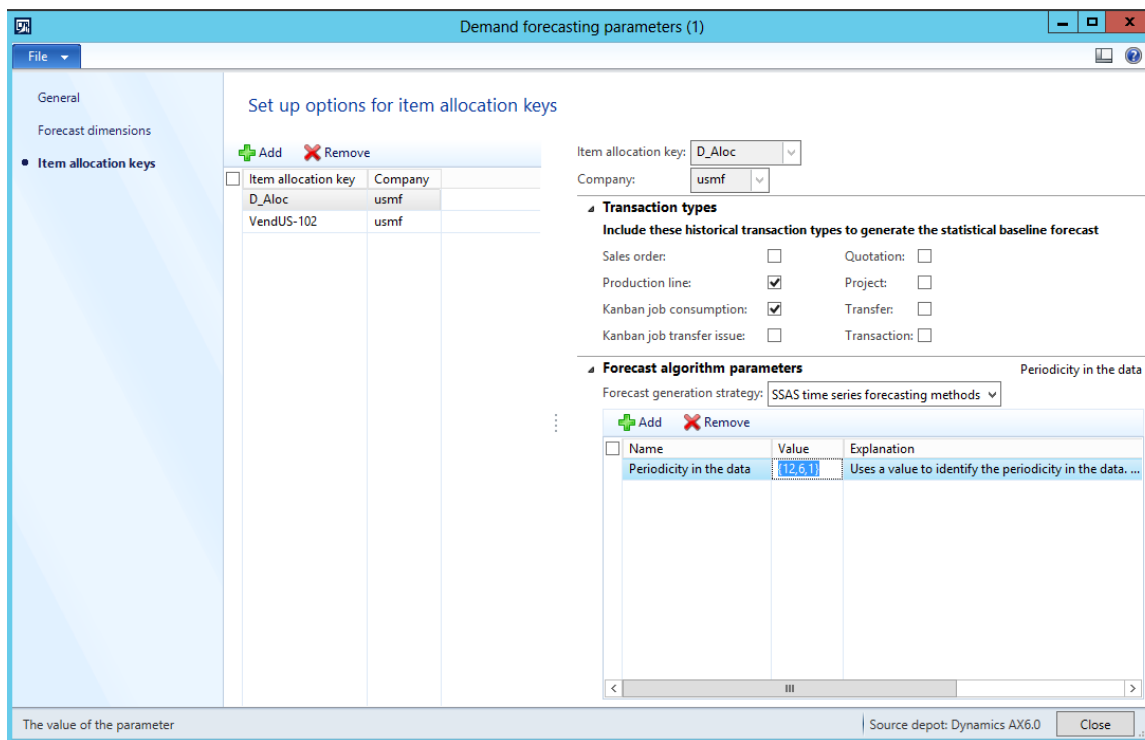
At any point, users can add forecast dimensions to the list of dimensions used for demand forecasting, or remove them, without having to redeploy the demand forecast and demand forecast accuracy cubes. However, manual adjustments are lost if you add or remove a forecast dimension.

Not all SKUs behave the same from a demand forecasting perspective. Similar SKUs can be grouped in one item allocation key, and parameters such as transaction types and Analysis Services forecast method settings can be set per item allocation key. To do this, click the **Item allocation keys** link in the left pane of the **Demand forecasting parameters** form.

Tip: When using SQL Server Enterprise Edition or SQL Server BI Edition, group items that have the same seasonality, trend, and base values for the forecast in the same item allocation key. This helps improve forecast accuracy, because items in the same allocation key are processed by the same mining model, and correlations between items are

determined. In SQL Server Standard Edition, there is just one mining model for all the items included in the forecast.

The following illustration shows the **Item allocation keys** area of the **Demand forecasting parameters** form.

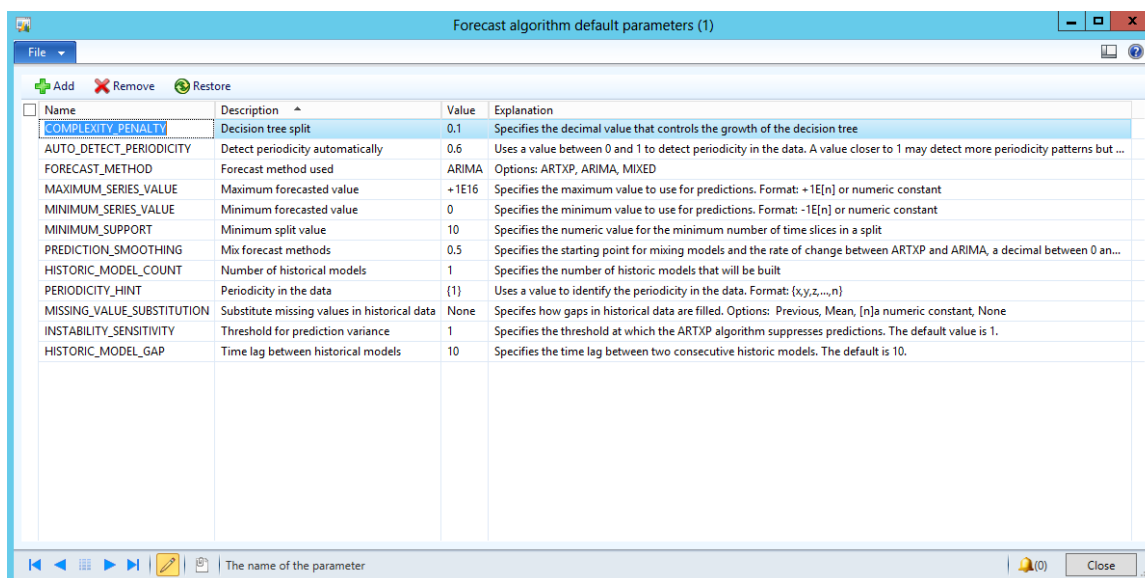


In the preceding illustration, to generate the demand forecast for the SKUs from item allocation key D_Aloc, the system will use all the default Analysis Services parameter values, except the periodicity hint value. This is overwritten to {12,6,1}, because these SKUs behave differently from a periodicity perspective.

SQL Server Analysis Services parameters setup

The set of parameters exposed by Analysis Services, and their default values, is accessed from Microsoft Dynamics AX, under **Master planning > Setup > Demand forecasting > Forecasting algorithm parameters**.

The following illustration shows the **Forecast algorithm default parameters** form.



This form is dedicated to power users who have knowledge of statistical forecasting and Analysis Services time series

forecasting. Power users can modify the descriptions of the parameters, so that they can be understood by users without a statistical forecast background. The default values from Analysis Services might not make sense for the data composition of all Microsoft Dynamics AX installations. Therefore, power users can overwrite the Analysis Services default values with values that make sense for the enterprise. Power users can also remove parameters that they decide make no sense for the Microsoft Dynamics AX installation. For example, if the ARIMA forecast method is used by an organization, the parameters that control the ARTxp method, such as *MINIMUM_SUPPORT*, can be deleted. At any point, the system can be restored to the Analysis Services default values by clicking **Restore**.

Users with no statistical forecasting knowledge can overwrite parameters based on their observations of the forecast accuracy behavior. Go to **Master planning > Setup > Demand forecasting > Demand forecasting parameters > Forecast algorithms parameters**, and then click **Add** to overwrite a parameter.

Not all products behave the same from a forecasting perspective. For example, different products have different seasonality. Therefore, the Analysis Services parameters can be further overwritten per item allocation key. Parameters that are not overwritten take the values defined under **Master planning > Setup > Demand forecasting > Demand forecasting parameters > General**. If no values are overwritten there, the parameters take the default values defined under **Master planning > Setup > Demand forecasting > Demand forecasting parameters > Forecast algorithms parameters**.

Generating a statistical baseline forecast

Demand forecasting generates a statistical, quantity-based baseline forecast. This is called a statistical forecast because it's calculated by using time series algorithms. It's called a baseline because it's meant to be a starting point that will be adjusted to take into account promotions, new customer additions, and so on.

To generate a demand forecast, go to **Master planning > Periodic > Demand forecasting > Generate statistical baseline forecast**.

The following illustration shows the **Create statistical baseline forecast** form.

The forecast bucket can be selected at forecast generation time. The available values are: **Day**, **Week**, and **Month**.

The number of buckets to generate a forecast for is set in the **Forecast horizon** field. In the preceding illustration, demand for 12 months will be generated.

When the forecast strategy is set to **Copy over historical demand**, the end of the historical horizon is ignored. AX 2012

R3 copies the number of buckets specified in the **Forecast horizon** field to the forecast demand, starting from the date set in the **Start** field under **Historical horizon**. By copying historical demand from a certain date forward, we let production planners make the plan for the next quarter in two ways:

- By copying the demand from the same quarter last year
- By copying the demand from the previous quarter

To prevent turbulence in the production plans, a certain number of forecast buckets can be frozen. This number is set in the **Freeze time fence** field. In the Excel workbook for the forecast, frozen buckets have a light yellow background to give a visual indication that those values should not be changed. However, there is no functionality to prevent users from modifying and publishing changes to the frozen forecast – the number of frozen buckets is just used to give a visual indication of values that should not be changed. The prohibition against overwriting forecasted quantities can be enforced only during import of adjusted demand forecasts.

The start date for the baseline demand forecast doesn't have to be the current date. To set a different start date, use the **Baseline forecast start date** field. For example, in June, users can generate a forecast for the next year. Because the forecast buckets between the end of historical demand and the start of the baseline are considered missing by Analysis Services, the predictions might not be very accurate. There are two ways to fill in this data, which will be discussed in the [Understanding the SQL Server Analysis Services forecasting methods](#) section.

The **Baseline forecast start date** field has to be set to the beginning of a forecast bucket (for example, in the United States, a Sunday if the forecasting bucket is the week). AX 2012 R3 automatically adjusts the **Baseline forecast start date** field to match the beginning of a forecast bucket.

The **Baseline forecast start date** field can be set to a date in the past. In other words, it is possible to generate a demand forecast in the past. This is useful, because it lets users tweak the Analysis Services parameters so that the statistical forecast generated in the past matches the actual historical demand. Users can then continue using these parameter settings to generate a statistical baseline forecast for the future.

Manual adjustments made in previous demand forecasting iterations can be automatically applied to the new baseline forecast if the **Transfer manual adjustments to the demand forecast** check box is selected. If the check box is cleared, the manual adjustments are not added to the baseline forecast inside the demand forecasting cube – but they are not deleted. Manual adjustments made to a forecast can be deleted only at forecast import time, by clearing the **Save the manual adjustments made to the baseline demand forecast** check box.

Manual adjustments are saved at import time. Therefore, if a user makes manual adjustments to the forecast, publishes the changes to the cube, but doesn't import the forecast back to Microsoft Dynamics AX, the changes are lost. For more information about manual adjustments and how they work, see the [Adjusted forecast import](#) section.

A demand forecast generation can have a name and comments to help users identify the forecast that has been generated. These values are visible in the Excel workbook used to visualize the demand forecast.

The path of the Excel workbook used to visualize and adjust the baseline demand forecast can be set in the **Excel demand forecast file path** field, but a default value is populated from the **Demand forecasting parameters** form. You can store the Excel file locally or on a SharePoint site.


The intercompany planning group, the item allocation keys, and other filters can be applied at forecast generation time. These can be used to improve performance or to split the data into manageable chunks. However, note that a demand forecast is not generated for the members of any item allocation key that is not associated with an intercompany planning group, even if the item allocation key is selected in the query.

Tip: Sometimes, users might receive errors while generating a demand forecast, or forecast generation is completed with no session log. This can happen because of leftover data in the query that was previously used for forecast generation. To fix this issue, click **Select** to open the **Query** form, click **Reset**, and then regenerate the baseline forecast.

After statistical baseline forecast has been generated, you can open the Excel file used to view the forecast by double-

clicking the message in the Infolog.

The following illustration shows the Infolog message.



Making manual adjustments to the baseline forecast

The statistical baseline forecast is calculated and stored inside the Microsoft Dynamics AX demand forecasting Analysis Services cube. Excel is used to connect to the demand forecasting cube so that users can visualize the forecast, and make comments and manual adjustments to it. Standard Excel pivot table and pivot chart functionality is used to visualize the historical and forecasted quantities.

The first worksheet of the Excel file, called Document information, contains a summary of the run.

The following illustration shows the Document information worksheet.

DemForecast3

FILE HOME INSERT PAGE LAYOUT FORMULAS DATA REVIEW VIEW Load Test Team

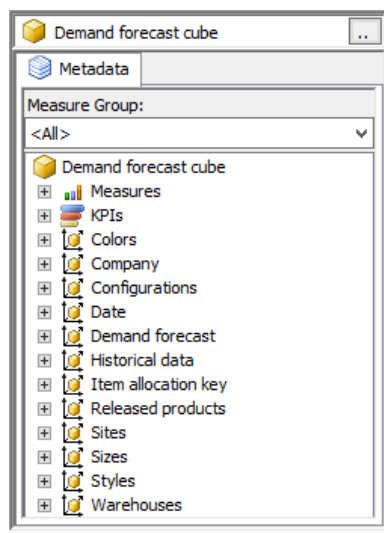
Clipboard Font Alignment Number

Generatio... : X ✓ fx 05-08-2013

	A	B	C	D	E	F	G
1	Statistical baseline forecast						
2	Generated date:	05-08-13					
3	Comment:	Demand forecast for 6 months, starting September 2013					
4	Forecast bucket:	Month					
5	Forecast horizon:	6					
6	Start of historical horizon:	01-01-12					
7	End of historical horizon:	30-11-12					
8	Baseline forecast start date:	01-09-13					
9							
10	Filters:						
11	Intercompany planning group	10					
12	Item allocation key	D_Aloc					
13							
14							
15							
16							

The second worksheet, Forecast table, contains the pivot table used to brows the cube data. The demand forecast cube has a fixed hierarchy – that is, a fixed set of dimensions data can be sliced upon.

The following illustration shows the demand forecast cube hierarchy.



The cube also has multiple measure groups: one for the actual demand quantities, one for the forecasted quantities, one for the changes made to the baseline forecast, and one that displays the demand forecast that is currently stored in the Microsoft Dynamics AX forecast models. There is also a measure that allows for an aggregate view of both historical and forecasted quantities. However, this view is read-only, so users can't make any adjustments to the forecast while using this view for the data.

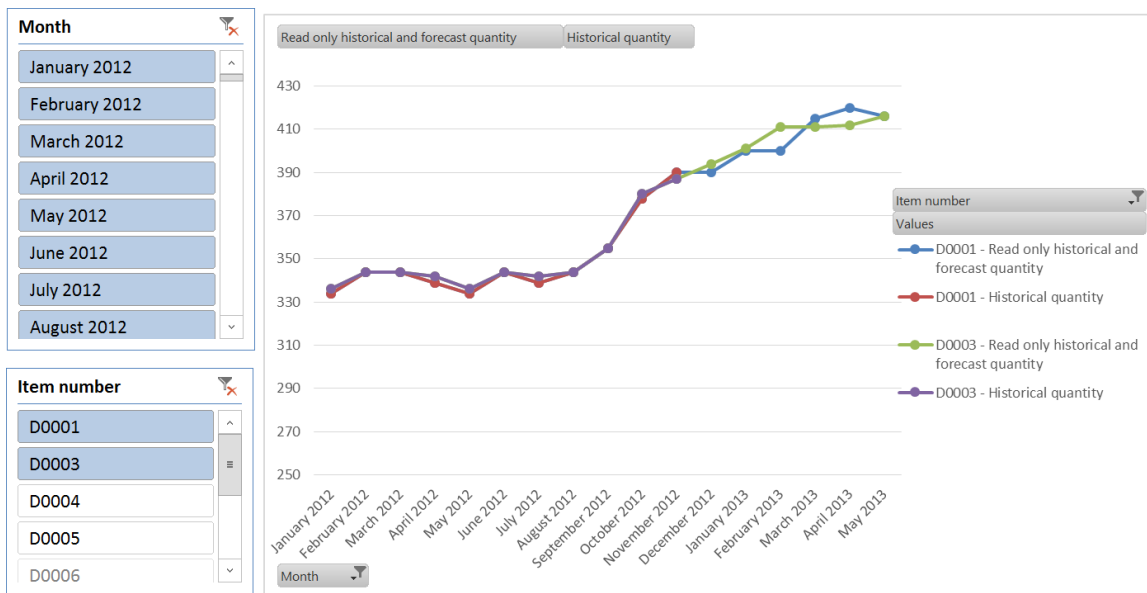
The following illustration shows the measure groups for the cube.

Row Labels	September 2012	October 2012	November 2012	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
D Alloc	1028	1089	1701	2030	1721	2853	1797	8023	1860	18284
D0001	355	378	390	450	422	440	455	467	481	2715
D0002			574	574	574	574	574	574	574	3444
D0003	355	380	387	397	407	420	426	434	443	2527
D0004	318	331	350	609	318	1419	342	6548	362	9598
Grand Total	1028	1089	1701	2030	1721	2853	1797	8023	1860	18284

In the Excel file, the cells representing the historical quantities have a grey background, the frozen periods have a light yellow background, and the forecasted quantities aren't highlighted in any way.

Users can also view a chart that shows the historical demand and the demand forecast. This lets them obtain a quick overview of how demand has been evolving, and whether demand forecast is off track. To get this view, add a pivot chart that is connected to the demand forecast cube to the Excel workbook. To produce the color-coded indication of history versus forecast, the chart needs to display both the historical quantity, and the read-only historical and forecast quantity measures.

The following illustration shows the chart of historical demand and the demand forecast for products D0001 and D0003.



Tip: Use the Insert Slicer functionality to slice data more easily.

Manual adjustments can be made to the forecast. These can be saved in the Excel workbook, but they need to be published back to the cube before they can be imported back into Microsoft Dynamics AX. To publish changes in Microsoft Excel 2013, go to **PivotTable Tools > Analyze > Calculations group > OLAP tools > What-If Analysis > Publish changes**. If manual adjustments have been made to the forecast but not published back to the cube, a red triangle is displayed in the lower-right corner of the affected cells.

To view the manual adjustments made to the baseline forecast, enable the **Adjusted quantity** measure on the pivot table.

The following illustration shows the pivot table after the **Adjusted quantity** measure is enabled.

	December 2012		January 2013		February 2013		March 2013		April 2013
Row Labels	Adjusted quantity	Forecasted quantity	Adjusted quantity	Forecasted quantity	Adjusted quantity	Forecasted quantity	Adjusted quantity	Forecasted quantity	Adjusted quantity
D Alloc	45.00	2030			1721	2853			1797
D0001	45.00	450			422	440			455
D0002		574			574	574			574
D0003		397			407	420			426
D0004		609			318	1419			342
Grand Total	45.00	2030			1721	2853			1797

PivotTable Fields

Show fields: (All)

- Current Ax forecast
 - Current Ax forecast
- Demand forecast
 - Adjusted quantity
 - Forecasted quantity
 - Read only historical and fore...

To view how far off the forecasted quantity is compared to actual demand that is already registered in AX 2012 R3, enable the **Actual demand** measure on the pivot table.

The following illustration shows the pivot table after the **Actual demand** measure is enabled.

	December 2012		January 2013		February 2013		March 2013		April 2013
Row Labels	Actual demand	Forecasted quantity	Actual demand	Forecasted quantity	Actual demand	Forecasted quantity	Actual demand	Forecasted quantity	Actual demand
D Alloc	1701	2030	1250	1721	117	2853	73	1797	
D0001	390	450	263	422	39	440	23	455	
D0002	574	574	222	574	36	574	26	574	
D0003	387	397	403	407	42	420	24	426	
D0004	350	609	362	318	36	1419	26	342	
Grand Total	1701	2030	1250	1721	117	2853	73	1797	

PivotTable Fields

Show fields: (All)

- Current Ax forecast
 - Current Ax forecast
- Demand forecast
 - Adjusted quantity
 - Forecasted quantity
 - Read only historical and fore...
- Historical data
 - Actual demand

To view a past forecast that has been used by AX 2012 R3 during planning and the current forecast that is registered in AX 2012 R3, enable the **Current Ax forecast** measure in the pivot table.

The following illustration shows the pivot table after the **Current Ax forecast** measure is enabled.

	December 2012		January 2013		February 2013		March 2013		April 2013
Row Labels	Forecasted quantity	Actual demand	Forecasted quantity	Actual demand	Forecasted quantity	Actual demand	Forecasted quantity	Actual demand	Forecasted quantity
D0001	450	263	450	263	422	39	440	23	455
D0002	574	222	574	222	574	36	574	26	574
D0003	397	403	397	403	407	42	420	24	426
D0004	609	362	609	362	318	36	1419	26	342
Grand Total	2030	1250	1721	117	2853	73	1797		

PivotTable Fields

Show fields: (All)

- Current Ax forecast
 - Current Ax forecast
- Demand forecast
 - Adjusted quantity
 - Forecasted quantity
 - Read only historical and fore...
- Historical data
 - Actual demand
- Colors
 - Colors
 - IsNotApplicable
- Company
 - Accounting currency

Drag fields between areas below:

FILTERS

COLUMNS

Released prod...

ROWS

Month

VALUES

Forecasted qu...

Actual demand

Current Ax for...

Handling the Excel issue with weighted allocation when cells that have a value of 0 are modified

Demand forecasting for AX 2012 R3 uses Analysis Services data mining to create the forecast and exposes the result of the prediction via an Excel document, in a pivot table. The predictions are stored in a cube with a write-enabled partition, which lets users perform what-if analysis. For the write-back, if the update happens on a higher level of aggregation in the Excel pivot table, weighted allocation is used to split the values. Weighted allocation lets users modify the aggregated value (for example, at the item allocation key, site, or company level), and this adjustment is proportionally propagated all the way to the SKU level.

For example, you want to change the aggregated value for site 1, warehouse 13, from 1103 to 1200.

The following example shows the pivot table before the change is made.

Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
1		1103	1111	1056	1075	990	1086	6421
13		1103	1111	1056	1075	990	1086	6421
D_Aloc		1103	1111	1056	1075	990	1086	6421
D0001		363	398	333	352	266	358	2070
D0003		394	401	411	411	412	416	2445
D0004		346	312	312	312	312	312	1906
4		426	426	426	426	426	426	2556
41		426	426	426	426	426	426	2556
D_Aloc		426	426	426	426	426	426	2556
D0005		426	426	426	426	426	426	2556
Grand Total		1529	1537	1482	1501	1416	1512	8977

After the aggregated value is changed to 1200, weighted allocation is triggered, and the adjustment is propagated all the way to the individual items.

The following illustration shows the pivot table after the change is made.

Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
1		1200	1111	1056	1075	990	1086	6518
13		1200	1111	1056	1075	990	1086	6518
D_Aloc		1200	1111	1056	1075	990	1086	6518
D0001		395	398	333	352	266	358	2102
D0003		429	401	411	411	412	416	2480
D0004		376	312	312	312	312	312	1936
4		426	426	426	426	426	426	2556
41		426	426	426	426	426	426	2556
D_Aloc		426	426	426	426	426	426	2556
D0005		426	426	426	426	426	426	2556
Grand Total		1626	1537	1482	1501	1416	1512	9074

However, there is an issue when there is a 0 (zero) value in the pivot table, either because the prediction is 0, or because a user modified the forecast to 0 and then performed a consecutive update of the same cell with a non-zero value. In this case, Analysis Services throws an error because of an exception while the cube is updated.


For example, the prediction for D0001 from the preceding example is changed to 0.

The following illustration shows the pivot table after this change is made.

Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
1		805	1111	1056	1075	990	1086	6123
13		805	1111	1056	1075	990	1086	6123
D_Aloc		805	1111	1056	1075	990	1086	6123
D0001		0	398	333	352	266	358	1707
D0003		429	401	411	411	412	416	2480
D0004		376	312	312	312	312	312	1936
4		426	426	426	426	426	426	2556
41		426	426	426	426	426	426	2556
D_Aloc		426	426	426	426	426	426	2556
D0005		426	426	426	426	426	426	2556
Grand Total		1231	1537	1482	1501	1416	1512	8679

The following illustration shows the error message that appears when the same prediction is later changed to 1.

Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
1		805	1111	1056	1075	990	1086	6123
13		805	1111	1056	1075	990	1086	6123
	D_Aloc	805	1111	1056	1075	990	1086	6123
	D0001	1	398	333	352	266	358	1707
	D0003	429	401	411	411	412	416	2480
	D0004	376	312	312	312	312	312	1936
4		426	426	426	426	426	426	2556
41		426	426	426	426	426	426	2556
	D_Aloc	426	426	426	426	426	426	2556
	D0005	426	426	426	426	426	426	2556
Grand Total		1231						



This happens because the following formula used by weighted allocation tries to divide by 0:

$$\text{Weight_Expression} = \text{<Leaf cell value>} / \text{<Existing value>}$$

A bug has been filed with the Analysis Services team to track this issue. However, until it is fixed, the only obvious workaround is to avoid the division by 0. Instead of 0, users should fill in 0.001, for example.

The following illustration shows the pivot table after the prediction for D0001 is changed to 0.001.

Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
1		740	1111	1056	1075	990	1086	6058
13		740	1111	1056	1075	990	1086	6058
	D_Aloc	740	1111	1056	1075	990	1086	6058
	D0001	0.001	398	333	352	266	358	1707
	D0003	394	401	411	411	412	416	2445
	D0004	346	312	312	312	312	312	1906
4		426	426	426	426	426	426	2556
41		426	426	426	426	426	426	2556
	D_Aloc	426	426	426	426	426	426	2556
	D0005	426	426	426	426	426	426	2556
Grand Total		1166	1537	1482	1501	1416	1512	8614

The following illustration shows the pivot table when the same prediction is later changed to 1.

Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
1		741	1111	1056	1075	990	1086	6059
13		741	1111	1056	1075	990	1086	6059
	D_Aloc	741	1111	1056	1075	990	1086	6059
	D0001	1	398	333	352	266	358	1708
	D0003	394	401	411	411	412	416	2445
	D0004	346	312	312	312	312	312	1906
4		426	426	426	426	426	426	2556
41		426	426	426	426	426	426	2556
	D_Aloc	426	426	426	426	426	426	2556
	D0005	426	426	426	426	426	426	2556
Grand Total		1167	1537	1482	1501	1416	1512	8615

Notice that there is no error message.

The preceding solution covers only the case of a manual change.

If a prediction from Analysis Services is 0, a manual adjustment on this prediction causes the same error as described previously. To prevent Analysis Services from predicting absolute 0, you can set the *MINIMUM_SERIES_VALUE* parameter to 0.001, for example. To do this, go to **Master planning > Setup > Demand forecasting > Forecasting algorithm parameters**.

Note: The *MINIMUM_SERIES_VALUE* parameter is available only in SQL Server Enterprise Edition.

If decimals make no sense for the items that you are forecasting, use the **Value field settings** option on the pivot table to make sure that no decimals are shown. In Excel, right-click in the pivot table, and then click **Value field settings**. Click **Number format**, and, for the **Number** category, make sure that **Decimals** is set to 0.

Making percentage adjustments to the forecast

For simulation purposes or to apply a business growth plan to your forecast, users might want to apply percentage adjustments to the demand forecast.

The next three scenarios explain how to achieve this. They all follow the same pattern: In Excel, apply the percentage change the first time this is needed; then, in subsequent forecasting iterations, apply the same change to any new forecast buckets that have been added.

Scenario 1: Every cycle, make an optimistic demand forecast that overall demand will increase by 10 percent across all companies and products

In the first scenario, there's no need to track manual adjustments made to the forecast, because the manual adjustment always applies a 10 percent increase. For this scenario, follow these steps.

1. Generate the forecast as usual, leaving the **Transfer manual adjustments to the demand forecast** check box cleared. In the comments, state that you intend to make an optimistic forecast of 10 percent overall growth.
2. Open the demand forecast Excel workbook.
3. Increase the **Grand Total** value in the lower right by 10 percent.

The following illustration shows the new **Grand Total** value.

Comments	Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
[15 Nov 2012] Increasing the forecasted values by 10%, based on predicted growth	USMF		1620	1652	1723	1762	1807	1852	10417
	D_Aloc		1620	1652	1723	1762	1807	1852	10417
	D0001		449	468	506	531	559	580	3092
	D0002		57	57	57	57	57	57	343
	D0003		436	454	487	502	519	543	2941
	D0004		210	205	204	204	204	204	1229
	D0005		469	469	469	469	469	469	2812
	Grand Total		1620	1652	1723	1762	1807	1852	10417

4. Publish the changes back to the cube. Manual adjustments other than the percentage change can be made in Excel and published to the cube, but these will not be tracked. Therefore, the next time the forecast is generated, you have to apply the same changes again.
5. Save the Excel file.
6. Import the forecast back into Microsoft Dynamics AX. The manual adjustments are not saved, so the **Import the manual adjustments made to the baseline demand forecast** check box does not have to be selected. For more information, see the next section.
7. During the next cycle, generate the baseline forecast again, as you did in step 1.
8. Open the Excel file. You will see the adjusted **Grand Total** value from last iteration. To see the new forecast values, you need to discard the changes previously made in Excel, by going to **PivotTable Tools > Analyze > Calculations > OLAP Tools > What-If analysis > Discard changes**.

The following illustration shows what your forecast now looks like.

Comments	Forecasted quantity	Column Labels	January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	Grand Total
	Row Labels								
	USMF		1262	1887	1324	1003	890	1576	7942
	D_Aloc		1262	1887	1324	1003	890	1576	7942
	D0001		448	182	462	205	462	205	1964
	D0002		1	107	58	116	1	54	337
	D0003		418	1064	1	1	1	753	2238
	D0004		15	53	207	195	66	104	640
	D0005		380	481	596	486	360	460	2763
	Grand Total		1262	1887	1324	1003	890	1576	7942

9. Make the same adjustment as in the previous iteration. (See step 3.)

In a slightly more complicated scenario, in addition to a manual increase of 10 percent on the overall forecast, you are making other manual adjustments to the forecast and want to keep these from cycle to cycle, so that the same changes do not have to be made over and over.

1. Generate the forecast as usual, selecting the **Transfer manual adjustments to the demand forecast** check box. In the comments, state that you intend to make an optimistic forecast of 10 percent overall growth.
2. Open the demand forecast Excel workbook.
3. Increase the **Grand Total** value in the lower right by 10 percent.
4. Make additional forecast adjustments. In this example, we are changing the value for D0001 for May to 600. The following illustration shows the adjustments.

Comments	Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
[15th Nov 2012] Increasing the forecasted values by 10% based on predicted growth	Row Labels								
	D_Aloc		1620	1652	1723	1762	1807	1873	10437
	D0001		449	468	506	531	559	600	3112
Changing the forecast for D0001, May 2013 based on promotion	D0002		57	57	57	57	57	57	343
	D0003		436	454	487	502	519	543	2941
	D0004		210	205	204	204	204	204	1229
	D0005		469	469	469	469	469	469	2812
	Grand Total		1620	1652	1723	1762	1807	1873	10437

5. Publish the changes back to the cube.
6. Save the Excel file.
7. Import the forecast back into Microsoft Dynamics AX. There is tracking of manual adjustments, so the **Import the manual adjustments made to the baseline demand forecast** check box should be selected. For more information, see the next section.
8. During the next cycle, generate the baseline forecast again, as you did in step 1, but modify the forecast start date and the historical horizon.
9. Notice that the 10-percent increase has been preserved, and that the adjusted value for May has been preserved. The following illustration shows what your forecast now looks like.

Comments	Forecasted quantity	Column Labels	January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	Grand Total
[15th Nov 2012] Increasing the forecasted values by 10% based on predicted growth	Row Labels								
	D_Aloc		1652	1723	1762	1807	1873	1576	10393
	D0001		468	506	531	559	600	205	2869
Changing the forecast for D0001, May 2013 based on promotion	D0002		57	57	57	57	57	54	340
	D0003		454	487	502	519	543	753	3259
	D0004		205	204	204	204	204	104	1123
	D0005		469	469	469	469	469	460	2803
	Grand Total		1652	1723	1762	1807	1873	1576	10393

However, you need to increase the value for June, because it's a new month that's not included in the last forecasting cycle.

10. Increase the **Grand Total** value for June by 10 percent. The following illustration shows the new **Grand Total** value.

Comments	Forecasted quantity	Column Labels						
[15th Nov 2012] Increasing the forecasted values by 10% based on predicted growth	Row Labels	January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	Grand Total
Changing the forecast for D0001, May 2013 based on promotion	D_Aloc	1652	1723	1762	1807	1873	1733	10550
	D0001	468	506	531	559	600	225	2889
	D0002	57	57	57	57	57	59	345
	D0003	454	487	502	519	543	828	3334
	D0004	205	204	204	204	204	114	1133
	D0005	469	469	469	469	469	506	2849
	Grand Total	1652	1723	1762	1807	1873	1733	10550

- Publish the changes back to the cube.
- Save the Excel file.
- Import the forecast back into Microsoft Dynamics AX. There is tracking of manual adjustments, so the **Import the manual adjustments made to the baseline demand forecast** check box should be selected. For more information, see the next section.

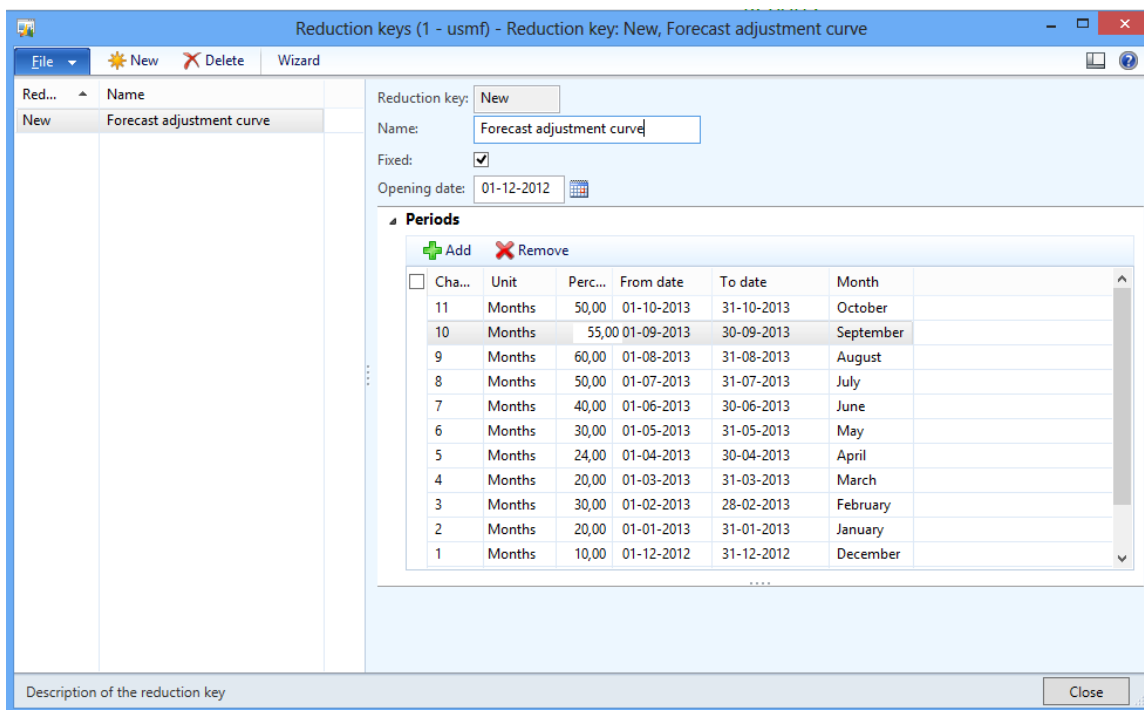
Scenario 2: Every cycle, make an optimistic demand forecast that overall demand will increase at a specified rate each month across all companies and products

In either the Excel forecast file or Microsoft Dynamics AX, the percentage curve that should be applied can be defined as a reduction key. You must then copy this reduction key to the Excel forecast file. After this, you can copy the **Grand Total** row to another location in the Excel file, apply the percentages, and copy either the values or the formula for calculating the percentage adjustment back to the pivot table.

Because it's the same change every month, dictated by the percentage, the manual adjustment is not saved between forecast generations.

- In Microsoft Dynamics AX, go to **Master planning > Setup > Coverage > Reduction keys**, and create a new reduction key. In this example, a forecast is generated starting December 2012, so a reduction key with the same start date is created.

The following illustration shows the reduction key.



- Press CTRL+T to export the reduction key to Excel. The following illustration shows the reduction key for export.

8. Save the Excel file.
9. Import the forecast back into Microsoft Dynamics AX. There is no tracking of manual adjustments, so it doesn't matter whether the **Import the manual adjustments made to the baseline demand forecast** check box is selected.
10. Start a new forecasting cycle – in this example, from January 2013. The following illustration shows the new forecasting cycle.

Comments	Forecasted quantity	Column Labels	January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	Grand Total
[15 Nov 2013] Adjusted forecast based on predicted growth curve	Row Labels								
	D_Aloc		1262	1887	1324	1003	890	1576	7942
	D0001		448	182	462	205	462	205	1964
	D0002		1	107	58	116	1	54	337
	D0003		418	1064	1	1	1	753	2238
	D0004		15	53	207	195	66	104	640
	D0005		380	481	596	486	360	460	2763
	Grand Total		1262	1887	1324	1003	890	1576	7942

11. Repeat steps 1 through 9, but make sure that the formulas are adjusted properly so that June, the newly added month, is also adjusted appropriately. Notice that the percentages are persisted in the Excel file. If you make any changes to the reduction key in Microsoft Dynamics AX, you have to manually change the values in Excel. The following illustration shows the new values.

Comments	Forecasted quantity	Column Labels	January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	Grand Total
[15 Dec 2013] Adjusted forecast based on predicted growth curve	Row Labels								
	D_Aloc		1514	2453	1589	1244	1157	2206	10163
	D0001		538	237	554	254	601	287	2470
	D0002		1	139	70	144	1	76	431
	D0003		502	1383	1	1	1	1054	2943
	D0004		18	69	248	242	86	146	808
	D0005		456	625	715	603	468	644	3511
	Grand Total		1514	2453	1589	1244	1157	2206	10163

Notice that, in step 5 of this example, the values forecasted for January, based on the forecast generation that happened in November 2012, are not the same as the values forecasted for January 2013 in the second iteration (step 10). This is because we also changed the historical horizon for the forecast to include December 2012 in the second iteration.

In a slightly more complicated scenario, in addition to a percentage adjustment on the overall forecast, you are making other manual adjustments to the forecast and want to keep these from cycle to cycle, so that the same changes do not have to be made over and over.

1. Generate the statistical baseline forecast, making sure that the **Transfer manual adjustments to the demand forecast** check box is selected.
2. In the Excel forecast document, fill in the percentages to apply.
3. Copy the **Grand Total** row to another place in the Excel workbook.
4. In the **Grand Total** field of the pivot table, write a formula that multiplies the old grand total value by the percentage, and copy it on the rest of the line.

The following illustration shows the new values.

Comments	Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
[15th Nov 2012] Adjust the demand forecast globally based on predicted growth	Row Labels								
	D_Aloc		1620	1802	2036	1922	2037	2189	11606
	D0001		449	510	598	579	630	685	3451
	D0002		57	62	68	62	64	68	382
	D0003		436	495	576	547	585	642	3281
	D0004		210	223	241	222	229	240	1366
	D0005		469	511	554	511	528	554	3126
	Grand Total		1620	1802	2036	1922	2037	2189	11606
			December	January	February	March	April	May	
	Copy of grand total		1473	1502	1566	1602	1643	1684	
	Adjustment percenta		1,1	1,2	1,3	1,2	1,24	1,3	

- Assume that, in May, a customer contract is signed that states that the customer will buy 500 of product D0001. Therefore, the forecast for May needs to be adjusted by at least 500. Obviously, this is a change that you will want to track and apply next time you generate the forecast. The following illustration shows the new values.

Comments	Forecasted quantity	Column Labels						
[15th Nov 2012] Adjust the demand forecast globally based on predicted growth	Row Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
	D_Aloc	1620	1802	2036	1922	2037	2504	11921
	D0001	449	510	598	579	630	1000	3766
	D0002	57	62	68	62	64	68	382
	D0003	436	495	576	547	585	642	3281
	D0004	210	223	241	222	229	240	1366
	D0005	469	511	554	511	528	554	3126
	Grand Total	1620	1802	2036	1922	2037	2504	11921
		December	January	February	March	April	May	
	Copy of grand total	1473	1502	1566	1602	1643	1684	
	Adjustment percenta	1,1	1,2	1,3	1,2	1,24	1,3	

- Publish the changes to the cube.
- Save the Excel file.
- Import the forecast back into Microsoft Dynamics AX. We are tracking changes and don't want to discard them, so make sure that the **Import the manual adjustments made to the baseline demand forecast** check box is selected.
- Start a new demand forecast iteration—in this example, from January 2013. The following illustration shows the new demand forecast iteration.

Comments	Forecasted quantity	Column Labels						
[15th Nov 2012] Adjust the demand forecast globally based on predicted growth	Row Labels	January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	Grand Total
	D_Aloc	1802	2036	1922	2037	2504	2206	12507
	D0001	510	598	579	630	1000	287	3604
	D0002	62	68	62	64	68	76	400
	D0003	495	576	547	585	642	1054	3900
	D0004	223	241	222	229	240	146	1301
	D0005	511	554	511	528	554	644	3302
	Grand Total	1802	2036	1922	2037	2504	2206	12507
		December	January	February	March	April	May	June
	Copy of grand total	1473	1502	1566	1602	1643	1684	1576
	Adjustment percenta	1,1	1,2	1,3	1,2	1,24	1,3	1,4

- Notice that the adjusted values are all there, and that the percentages are still kept in the Excel file. You just need to adjust June 2013 for the predicted growth.
- Publish the changes.
- Save the Excel demand forecast file.
- Import the forecast back into Microsoft Dynamics AX. We are tracking changes and don't want to discard them, so make sure that the **Import the manual adjustments made to the baseline demand forecast** check box is selected.

Scenario 3 – Phase-out: For a product that is reaching the end of its life, decrease the forecast by a certain percentage each month

Although this scenario refers to phasing out products, the same procedure can be applied when a product had good sales in the past (for example, because of promotions or a big customer orders), but you now have to level down the forecast. Basically, these steps can be applied whenever an increase or decrease in demand needs to be made for a specific SKU.

- Generate the statistical baseline forecast, making sure that the **Transfer manual adjustments to the demand forecast** check box is selected.
- Outside Microsoft Dynamics AX or Excel, define a plan to phase out a product. In this example, products D0001 and D0002 need to be phased out.
- In Excel, save the phase-out rules. The following illustration shows the phase-out rules.

Comments	Forecasted quantity	Column Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
[15 Nov 2013] Starting to phase out D0001 and D0002	D_Aloc		1407	1392	1889	1369	1339	1318	8714
	D0001		367	340	422	290	254	211	1884
	D0002		27	27	38	12	2	2	108
	D0003		396	413	581	456	472	494	2812
	D0004		191	186	243	185	185	185	1175
	D0005		426	426	605	426	426	426	2735
	Grand Total		1407	1392	1889	1369	1339	1318	8714
			December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	
	Product D0001 forecast copy		408	425	603	483	508	527	
	Product D0001 reduction percentages		0,9	0,8	0,7	0,6	0,5	0,4	
	Product D0002 forecast copy		52	52	68	52	52	52	
	Product D0002 forecast reduction quantity		25	25	30	40	50	50	

4. Save the Excel file.
5. Publish the changes.
6. Import the forecast back into Microsoft Dynamics AX. We are tracking changes and don't want to discard them, so make sure that the **Import the manual adjustments made to the baseline demand forecast** check box is selected.
7. Start a new forecasting cycle, from January 2013.
8. Open the Excel forecast file. Notice that the phase-out rules are still in the Excel workbook. Also notice that the forecast for D0001 is already trending to lower values, so we don't need to apply any more percentage adjustment for D0001. Forecast adjustment is needed only for D0002. The following illustration shows the forecast.

Comments	Forecasted quantity	Column Labels	January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	Grand Total
[15 Nov 2013] Starting to phase out D0001 and D0002	D_Aloc		1180	2058	1106	938	640	1523	7445
	D0001		340	422	290	254	211	205	1722
	D0002		27	38	12	2	2	1	82
	D0003		418	1064	1	1	1	753	2238
	D0004		15	53	207	195	66	104	640
	D0005		380	481	596	486	360	460	2763
	Grand Total		1180	2058	1106	938	640	1523	7445
			January 2013	February 2013	March 2013	April 2013	May 2013	June 2013	
	Product D0001 forecast copy		425	603	483	508	527	205	
	Product D0001 reduction percentages		0,8	0,7	0,6	0,5	0,4		
	Product D0002 forecast copy		52	68	52	52	52	54	
	Product D0002 forecast reduction quantity		25	30	40	50	50	53	

9. Save the file.
10. Publish the changes.
11. Import the forecast back into Microsoft Dynamics AX. We are tracking changes and don't want to discard them, so make sure that the **Import the manual adjustments made to the baseline demand forecast** check box is selected.

Scenario 4 – Phase-in: Create a forecast for a new product by copying the forecast from another product and adjusting it by a certain percentage

New products don't have historical demand, so Analysis Services can't generate a forecast for them. In addition, AX 2012 R3 can't copy the historical forecast as demand.

You can include these items in the forecast, but the only way to obtain a forecasted quantity for them is by copying and making percentage adjustments to demand forecast or historical demand data of other products. As a prerequisite, the new products need to be created in AX 2012 R3 and released to the corresponding companies. They also need to have a series of default values set up, such as a unit of measure, site, and warehouse. They also need to be part of an item allocation key, so they can be taken into account for forecast generation.

In the following example, we are creating a demand forecast for a new product, assuming that its demand behaves the same as the demand for another product, and adjusting the forecast by a certain percentage. You can consider this a sort of phase-in rule that is saved inside the Excel forecast file.

1. Create a new product:
 - **Product type** = Item
 - **Product subtype** = Product
 - **Product number** = NewD0008
 - **Item model group** = FIFO
 - **Item group** = Audio
 - **Storage dimension group** = SiteWH
 - **Tracking dimension group** = None
 - **Unit of measure** = ea (for Inventory, Sales, and Purchase)
2. Under **Demand forecast parameters > Forecast dimensions**, make sure that only Stock keeping unit is used as a forecast dimension.

Site is a mandatory forecast dimension. The new product has no demand history, so a site value cannot be determined for the forecast by looking into historical transactions. Therefore, to include the new product for demand forecasting, make sure that a default inventory site is specified in the **Default order settings** form, under **Product information management > Common > Released products > Plan > Order settings**. If Warehouse is also used as a forecast dimension, the new product must also have a default inventory warehouse.
3. Create a new item allocation key to group all the new products. This makes it easier to identify the new products:
 - **Item allocation key** = NewProducts
 - **Name** = NewProducts
4. Add the new product, NewD0008, to the new item allocation key, NewProducts.
5. Generate the demand forecast:
 - **Historical horizon start** = 01 Jan 2012
 - **Historical horizon end** = 30 Nov 2012
 - **Baseline forecast start date** = 01 Dec 2012
 - **Forecast bucket** = Month
 - **Forecast horizon** = 6

The historical date interval doesn't really matter when you generate a forecast for new products.

When you create a forecast for new products, it doesn't matter what **Forecast generation strategy** value is selected in the **Demand forecast parameters** form.

It's easy to recognize new products, because they have a flat forecast quantity, with no increase or decrease in the forecasted quantities. This flat forecast quantity is equal to the value provided by the *MINIMUM_SERIES_VALUE* parameter.

The following illustration shows the forecast for the new product, NewD0008.

Forecasted quantity	Column Labels						
Row Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
D_Aloc	1242	1163	1108	1127	1042	1138	6820
D0001	450	398	333	352	266	358	2157
D0002	52	52	52	52	52	52	312
D0003	394	401	411	411	412	416	2445
D0004	346	312	312	312	312	312	1906
NewProduct	1	1	1	1	1	1	6
NewD0008	1	1	1	1	1	1	6

6. Assume that item NewD0008 behaves like item D0001; however, because it's new, we need to adjust its demand by certain percentages. In Excel, create a phase-in rule. The following illustration shows the new phase-in rule.

	Dec-12	Jan-13	Feb-13	Mar-13	Apr-13	May-13
Product D0001 copy forecast	450	398	333	352	266	358
Percentage reduction rule for NewD0008	0.4	0.5	0.6	0.7	0.8	0.9
Forecast for NewD0008	180	199	200	246	213	322

7. Copy the forecast for the new product to the pivot table. The following illustration shows the new rule in the pivot table.

Forecasted quantity	Column Labels						
Row Labels	December 2012	January 2013	February 2013	March 2013	April 2013	May 2013	Grand Total
D_Aloc	1242	1163	1108	1127	1042	1138	6820
D0001	450	398	333	352	266	358	2157
D0002	52	52	52	52	52	52	312
D0003	394	401	411	411	412	416	2445
D0004	346	312	312	312	312	312	1906
NewProduct	180	199	200	246	213	322	1360
NewD0008	180	199	200	246	213	322	1360
Grand Total	1422	1362	1307.8	1373.4	1254.8	1460.2	8180.2
	Dec-12	Jan-13	Feb-13	Mar-13	Apr-13	May-13	
Product D0001 copy forecast	450	398	333	352	266	358	
Percentage reduction rule for NewD0008	0.4	0.5	0.6	0.7	0.8	0.9	
Forecast for NewD0008	180	199	200	246	213	322	

8. Publish the changes.
 9. Save the Excel file, and exit Excel.
 10. Import the forecast back into Microsoft Dynamics AX. Most likely, you don't want to discard the adjustments you made to the new product, so make sure that the **Save the manual adjustments made to the baseline demand forecast** check box is selected.
 11. You have to repeat steps 1 through 10 until the product has reached its maturity, when you can rely on the product's historical demand to create the forecast. After that, you can just let Analysis Services create the forecast.

Importing an adjusted forecast

After manual adjustments have been made to the forecast in Excel and published to the demand forecasting cube, forecast lines can be imported back into the Microsoft Dynamics AX transactional database, so that the forecast can be consumed by Material requirements planning (MRP). The demand forecast isn't saved in the Excel workbook but in the SQL Server data cube, and that's where it is imported from.

To import an adjusted demand forecast back into Microsoft Dynamics AX, go to **Master planning > Periodic > Demand forecasting > Import forecast**.

The following illustration shows the **Import forecast** form.

Import forecast (1)

Import adjusted demand forecast

General | Details | Batch

Import settings

Import forecast period:

Start: 01-Dec-12 End: 31-May-13 Restore defaults 01-Dec-2012 to 31-May

Save the manual adjustments made to the baseline demand forecast

Select the companies and the corresponding forecast models in which the adjusted demand forecast will be imported

+ Add - Remove

<input type="checkbox"/>	Company	Forecast model
<input type="checkbox"/>	usmf	CurrentF

OK Cancel

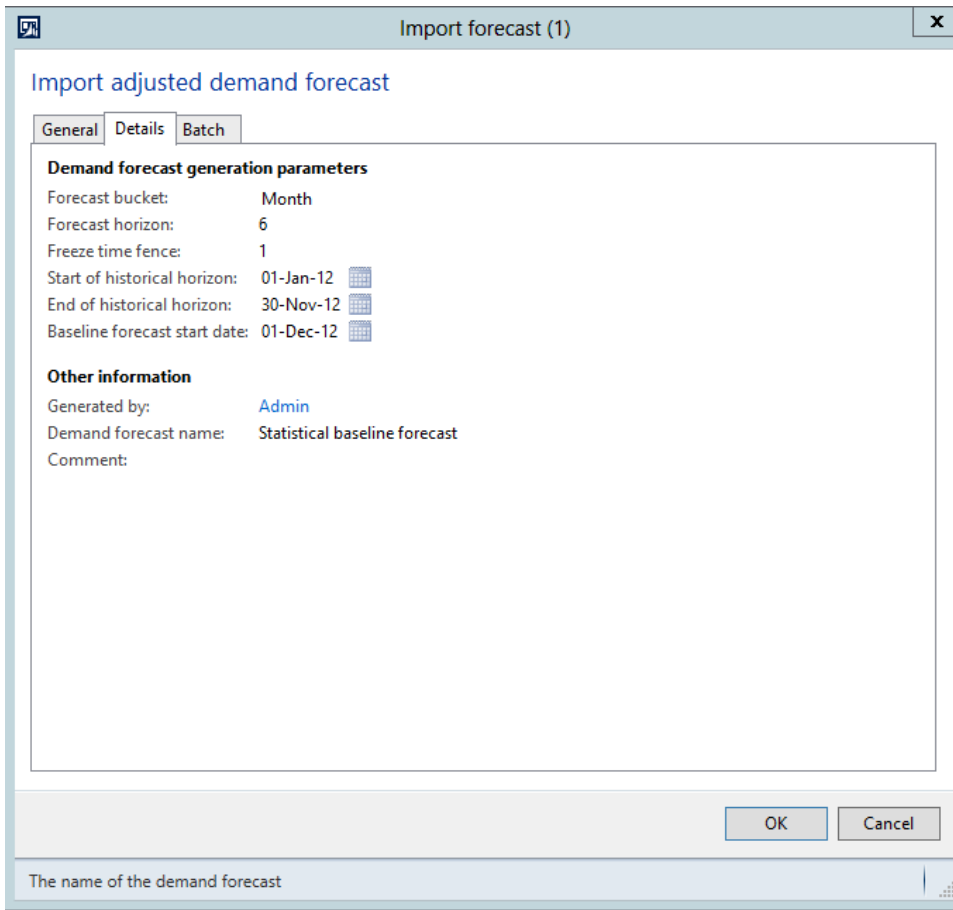
Delete the selected record

For the items included in the forecast, the import function deletes all existing forecast lines from the AX 2012 R3 database and creates new ones that correspond to the records found in the cube.

Not all forecast data from the demand forecasting cube has to be imported back into AX 2012 R3. Users can choose the start and end dates of the period that the forecast can be imported into. This lets users freeze certain buckets. The start and end dates need to correspond to start and end dates of the bucket that the forecast is generated in. AX 2012 R3 makes sure that this restriction is enforced and adjusts the dates automatically.

Users can view details about the forecast that is currently saved in the cube on the **Details** tab.

The following illustration shows the **Details** tab.



At import time, users can select the companies and the forecast models to import the forecast into. The grid is prefilled with all the companies that have forecast demand in the cube. For each company, the forecast model corresponding to the current forecast plan set up in the master planning parameters is also prefilled, but it can be changed to any forecast model belonging to the specific company. If a selected company has no forecast demand data in the cube, a warning is thrown at import time.

It is very important to understand how the **Save the manual adjustments made to the baseline demand forecast** check box works. If a user has made manual adjustments to the statistical baseline forecast and published them to the cube, the adjusted values are imported into AX 2012 R3, even if the check box is cleared. However, the changes are discarded after the import. Therefore, the next time a forecast is generated, even if **Transfer manual adjustments to the demand forecast** is selected, the forecast that is created is only the statistical one, without any manual overrides. So, in a way, the **Save the manual adjustments made to the baseline demand forecast** check box is a mechanism that lets users keep or discard all their manual changes.

Monitoring forecast accuracy

Forecast accuracy is a measurement of forecast usefulness. This is often defined as the mean absolute deviation (MAD), or the average difference between the forecasted value and the actual value.

AX 2012 R3 offers two forecast accuracy KPIs, MAD and MAPE (mean absolute percentage error).

The following illustrations show how these two KPIs are defined.

• Mean Absolute Deviation

$$MAD = \frac{1}{n} \sum_{i=1}^n |A_i - F_i|$$

Where :

MAD : Mean Absolute Deviation

A_i : The actual value in time period i

F_i : The forecast value in time period i

- Mean Absolute Percentage Error

$$MAPE = \frac{1}{n} \left[\sum_{i=1}^n \left| \frac{A_i - F_i}{A_i} \right| \right] \times 100$$

Where :

MAPE : Mean Square Error

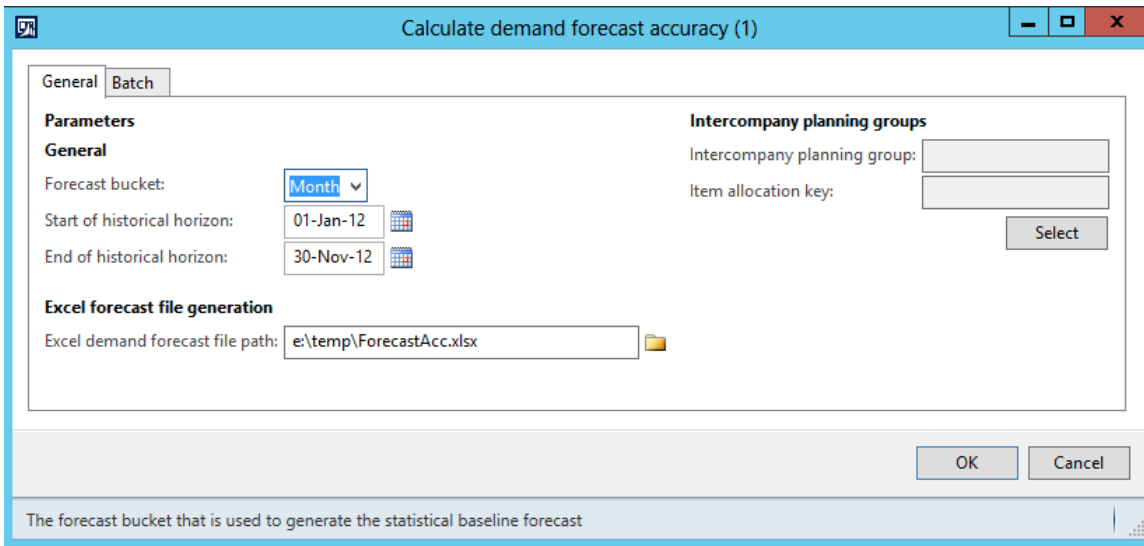
A_i : The actual value in time period i

F_i : The forecast value in time period i

The KPI values, together with the actual demand and forecast, are stored inside the demand forecast accuracy data cube. To visualize the data from this cube, use the same Excel workbook that you use to visualize the baseline forecast.

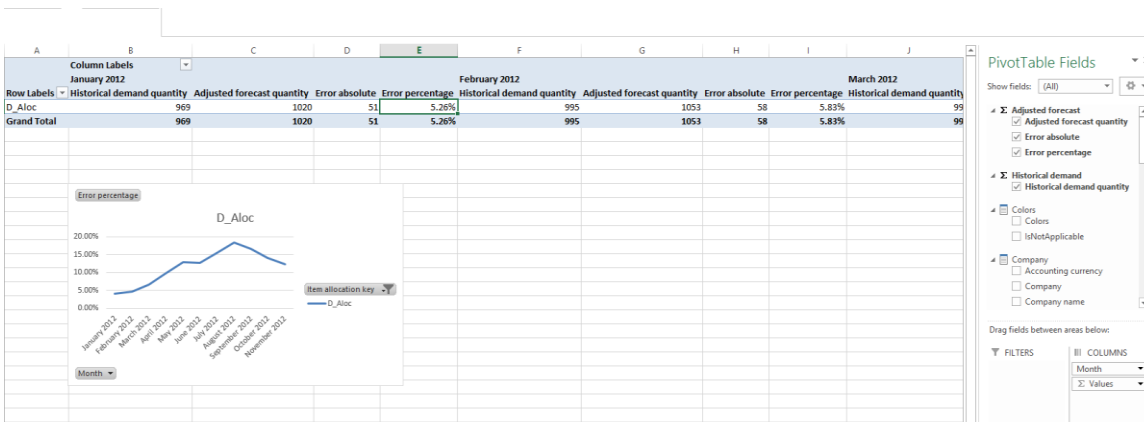
To trigger the KPI calculation, go to **Master planning > Periodic > Demand forecasting > Calculate demand forecast accuracy**.

The following illustration shows the **Calculate demand forecast accuracy** form.



The forecast accuracy measurements can be viewed in an Excel workbook. MAD is represented by the **Error absolute** measure and MAPE by the **Error percentage** measure.

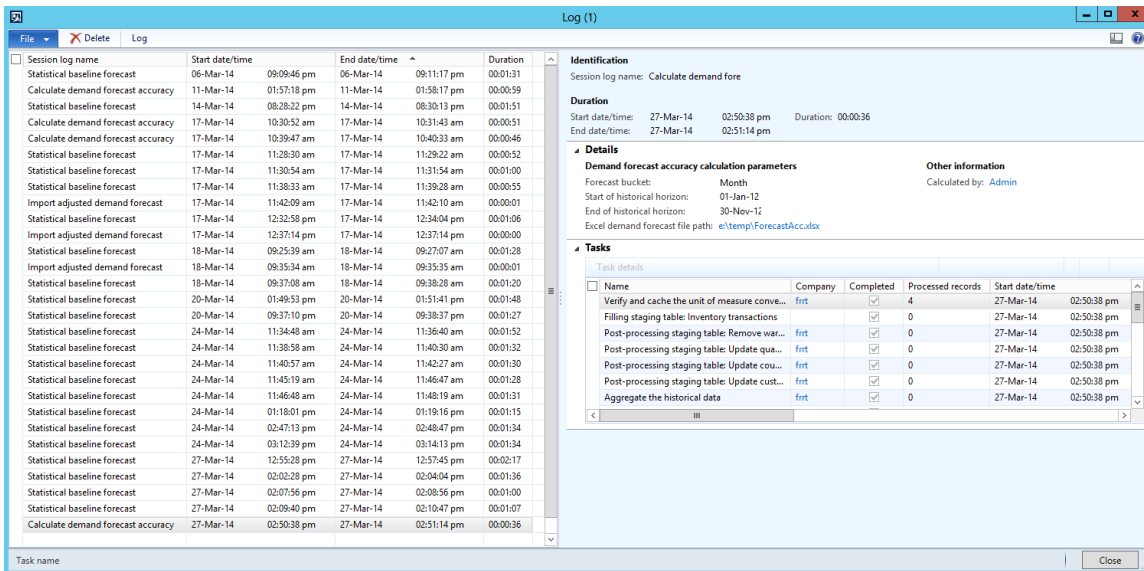
The following illustration shows these measures in the Excel workbook.



Monitoring forecast generation and import

Users can monitor forecast generation, forecast import, and forecast accuracy calculation tasks from the **Log** form, under **Master planning > Periodic > Demand forecasting > Session log**.

The following illustration shows the **Log** form.



The form shows information such as the start and end times, the duration of the task, the user who triggered the task, and the error log of the task.

Understanding the Analysis Services forecasting methods

Analysis Services offers advanced forecast models, which let users quickly generate highly accurate forecasts. Some customers have found these forecast models difficult to understand. Therefore, this section of the white paper aims to provide more insight into how these forecast models work, to help more customers understand and gain confidence in them.

AX 2012 R3 DemoData

This section of the white paper uses the AX 2012 R3 DemoData dataset to explain the Analysis Services time series forecasting methods and the parameters that control them. Therefore, it is important to become accustomed to this dataset.

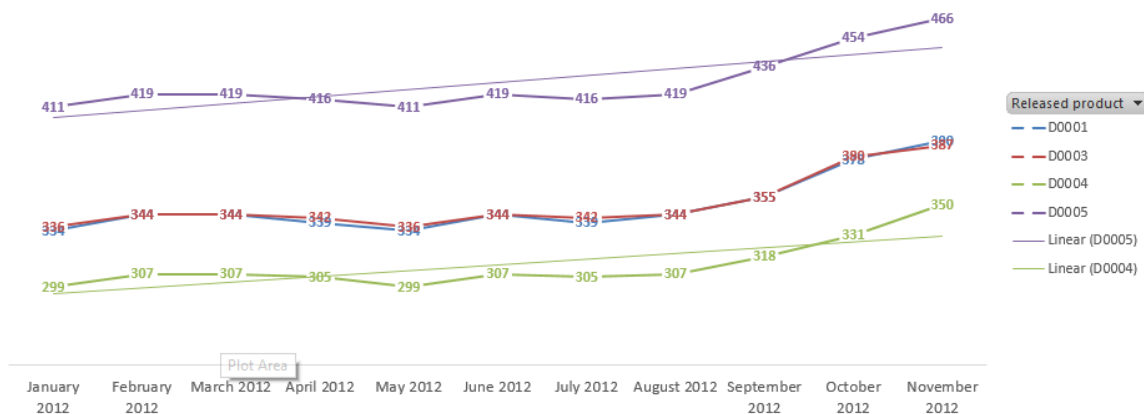
The examples in this section are based on sales history only, from company USMF. We will generate forecast only for the finished goods that the company sells.

Before a released product can be taken into account for demand forecasting, it needs to be part of an item allocation key. In company USMF, there is an item allocation key called D_Alloc, which contains the following released products: D0001, D0003, D0004, and D0005.

These products have both sales history and demand forecast lines between January 1, 2012, and November 30, 2012. For an overview of the sales trend per product, generate forecast accuracy for the period between January 1, 2012, and November 30, 2012. Add a pivot chart in the Excel workbook, and enable the **Historical demand quantity** measure in the chart.

The following illustration shows the overview of the sales trend for each released product. You can use this illustration instead of generating forecast accuracy for yourself.

Historical demand quantity



Because we don't have a lot of history, we will use as few forecast dimensions as possible when aggregating the data: Company, Item allocation key, and Stock keeping unit.

Demand forecasting approaches

There are two types of demand forecasting approaches: subjective/qualitative and objective/quantitative.

The subjective forecasting approach is primarily based on the expert's knowledge and opinions about market conditions, customers, and competition. This approach incorporates strategies such as the Delphi method, historical life cycles of similar products, and market research. No statistics are involved in creating the forecast, just human knowledge.

The objective forecasting approach is based on mathematical formulations (statistics). Essentially, the quantitative forecasting approach takes the amounts or quantities sold in the past to forecast how much will be sold in the near future. Therefore, it's obvious that this approach works only if there is enough relevant numerical historical data.

There are two types of objective forecasting approaches:

- **Time series models** – These are used to forecast continuous values over time. The only factor they consider when creating the forecast is time; mathematically speaking, time is the only variable the statistical method uses to analyze trends, seasonality, and cyclical and random factors. There are two types of time series models:
 - **Smoothing models** – Examples include moving average (simple and weighted), and single and double exponential smoothing.
 - **Decomposition models** – These can be additive or multiplicative.
- **Casual models** – These use more than just time as a variable. For example, weather forecasts, holidays, a soccer team's schedule, and correlations between these variables are used when the forecast is created.

AX 2012 R3 offers support for both the subjective and objective forecasting approaches.

Customers who prefer to use subjective forecast approaches can create a baseline by copying over historical demand as a forecast. To do this, select **Copy over historical demand** in the **Forecast generation strategy** field in the **Demand forecast parameters** form.

Customers who prefer to use objective forecasting approaches can use the time series Analysis Services forecast models, which will be explained in the next section of the white paper. To do this, select **SSAS time series forecasting methods demand** in the **Forecast generation strategy** field in the **Demand forecast parameters** form.


Introduction to Analysis Services time series forecasting

This section consolidates information that is available on TechNet and at <http://www.sqlserverdatamining.com>, or on

MSDN. For even more details, see the sources in the [References](#) section.

Analysis Services offers only time series forecasting methods. A *time series* is a set of values observed over a period of time, typically at regular intervals. Monthly sales quantities are a common example. Time series data is often presented in a graphical format, with the time interval along the x-axis of a chart and the values along the y-axis.

The following illustration shows an example of a chart of time series data.



As explained in the article [Understanding Time Series Forecasting Concepts](#), time series data has four major characteristics that affect how a value changes from one period to the next and how future values are forecasted:

- **Base level** – The base level is typically defined as the average value of the series. In some forecasting models, the base level is defined as the starting value of the series data.
- **Trend** – A trend is typically defined as how the series changes from one period to the next. In the preceding chart, the trends are calculated for products D0004 and D0005.
- **Seasonality** – Seasonality is a repetitive pattern found in the data. Certain values tend to rise and fall based on certain time periods, such as the day of the week or month of the year. For example, retail sales often spike during the Christmas season.
- **Noise** – Noise consists of random variations and irregular movements in the data.

If you can identify a trend, apply that trend to the base level, and account for any seasonality that might exist in the data, the following forecasting model can be used to predict future values:

$$\text{Forecasted value} = \text{Base level} + \text{Trend} + \text{Seasonality}$$

For more information about base levels, trends, and seasonality, see [Understanding Time Series Forecasting Concepts](#).

The preceding equation basically represents a mining model. The model is created by applying an algorithm to data. However, what is applied is more than an algorithm: it's a set of data, statistics, and patterns that can be applied to new data to generate predictions and make inferences about relationships.

Analysis Services time series methods can predict trends based only on the original time series dataset that is used to create the mining model. The patterns that the algorithms look for are inherent in the data itself. For example, given a series of stock prices, if we want to predict what the next week's stock prices will be, the algorithms look into the existing series, find the patterns, and project them forward to determine next week's stock prices. It is unlikely that looking at the historical stock prices alone will give an accurate forecast, especially in the long term, because stock prices depend on various factors, such as the US Federal Bank interest rates. These factors are not inherent in the existing stock price data. To achieve better accuracy, Analysis Services can do cross-prediction on time series data that is somehow related. However, there is no way to bring outside knowledge into the forecast algorithms.

The Analysis Services time series algorithm is essentially two algorithms blended into one: ARTxp and ARIMA. This blend is used, because it can be difficult for users who have no statistical background to understand which method to choose based on the forecast horizon they want. Basically, it's not easy to identify the forecast horizon before which ARTxp should be used and the horizon after which ARIMA should be used. For the blended algorithm, the mining model is trained once with the ARTxp algorithm and once with the ARIMA algorithm. Then SQL Server tries to determine the point to which the ARTxp results are copied to the forecast; after that, the ARIMA results are copied to the forecast. SQL Server Enterprise Edition includes a parameter that lets users control the blend. This parameter will be explained in detail later.

The ARTxp forecasting algorithm

The ARTxp algorithm comes from Microsoft Research and is based on Microsoft Decision Trees; therefore, it's unique to Microsoft. It was first time introduced in Microsoft SQL Server 2005.

The ARTxp algorithm supports cross-prediction, which is what the "xp" stands for.

The ARTxp algorithm is optimized to predict the next likely values (the next day's value, the next week's value, or the next month's value), based on the given time series data. We do not recommend that users who want to predict more than five forecast buckets in the future use only the ARTxp method, because the algorithm tends to become very noisy.

The ARIMA forecasting algorithm

The ARIMA forecasting algorithm was added to SQL Server 2008 to smooth out long-term predictions. Essentially, it's an implementation of the Box-Jenkins forecasting method.

The ARIMA algorithm doesn't support cross-time series prediction.

Demand forecasting parameters

The demand forecasting parameters control the behavior, performance, and accuracy of the mining model.

Important! Although the parameters are presented separately, avoid thinking about each parameter in isolation when you are trying to understand its effect on the forecast.

AUTO_DETECT_PERIODICITY

This parameter applies to both the ARTxp and ARIMA methods. Analysis Services used Fast Fourier transformations to detect seasonality before the training of the mining model.

The parameter takes a value between 0 and 1. If the value is closer to 0, periodicity is detected only for strongly periodic data. A value closer to 1 favors the discovery of many patterns that are almost periodic and the automatic generation of periodicity hints.

PERIODICITY_HINT

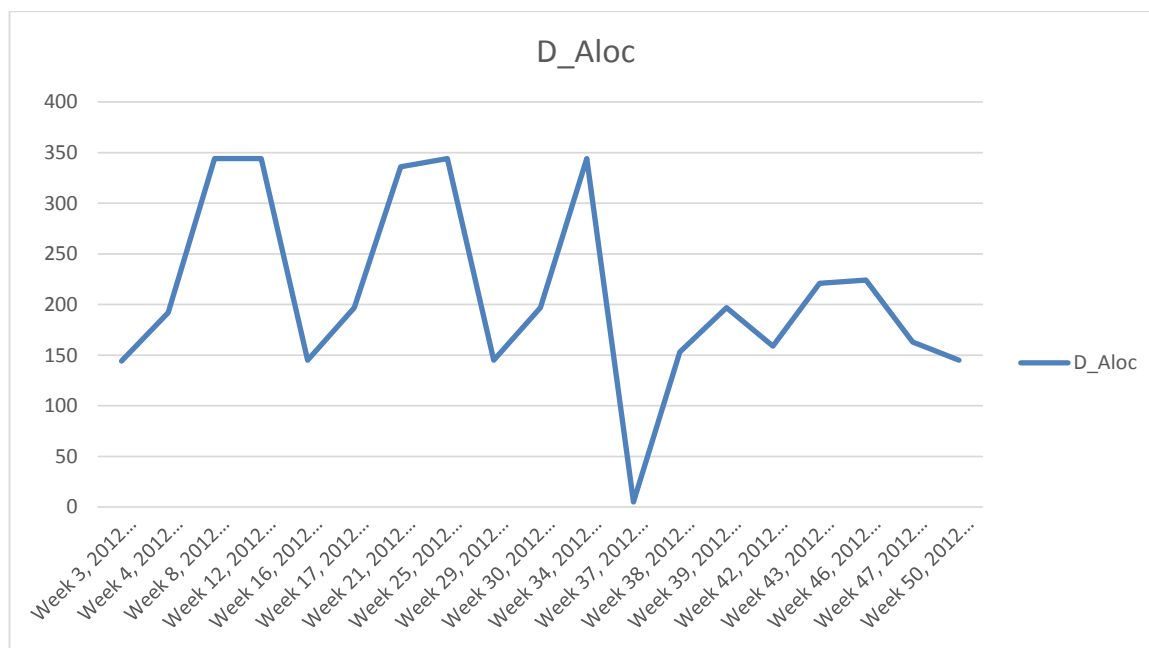
This parameter applies to both the ARTxp and ARIMA methods. Both are very sensitive to periodicity hints, so providing the wrong hint can adversely affect the results.

If you know the periodicity of the time series data, we highly recommend that you give the Analysis Services time series forecasting methods hints about what periodicity the mining model should use. For example, the sales of a company vary from year to year. Most commonly within a year, Month is used as the forecast bucket. Therefore, the periodicity of the data is {12} (note that this is exactly the format that users should enter data in the **Forecast algorithms parameters form** or the **Demand forecast parameters form**). But within a year, there are other periodicities: the sales vary from quarter to quarter, and from month to month within a given quarter, based on the change of the seasons, for example. Therefore, you can give more periodicity hints, such as {12, 3, 1}, which hints the models to detect patterns for a year, for a quarter within the year, and for a month within the quarter.

By default, the parameter value is {1}, which essentially hints the models that they should search for periodicity inside every bucket. In many cases, this is basically equivalent to no periodicity hint.

For an example of the effect of this parameter, let's look at the sales history of item D0003 in DemoData, in weekly buckets. We are using weekly buckets, because we already have periodicity at the week level for item D0003 in DemoData: every 12 weeks, some sort of cycle is completed, and within the 12 weeks, a decrease in demand occurs every six weeks.

The following illustration shows the sales history of item D0003.




Note: In all the examples that follow, because we are forecasting more than five data points, we will use ARIMA as the only method for forecasting (*FORECAST_METHOD=ARIMA*). This is done for two reasons: we don't want to get the blending parameters involved at this stage, and, more importantly, ARTxp tends to become very noisy when more than five data points are predicted.

Generate a 12-week forecast by setting the following values in the **Generate statistical baseline forecast** form:


- **Historical horizon start date** = 01 January 2012
- **Historical horizon end date** = 01 December 2012
- **Baseline forecast start date** = 02 December 2012
- **Forecast bucket** = Week
- **Forecast horizon** = 12

The following illustration shows a 12-week forecast as predicted by the ARIMA model, with very limited periodicity detection (*AUTO_DETECT_PERIODICITY=0* and *PERIODICITY_HINT={1}*).



This doesn't look very good, does it?


The following illustration shows the forecast if we let the ARIMA model identify the periodicity on its own (*AUTO_DETECT_PERIODICITY=1* and *PERIODICITY_HINT={1}*).



Here, the model identified the periodicity quite well, but the predicted values appear to be off the chart.

In DemoData, because we created the sales so that they vary year by year, we have at least {52}, the number of weeks in a year, as the periodicity hint. However, if we look at the data carefully, we can see that the same demand pattern repeats itself every 12 weeks: during each 12-week period, sales increase and then decrease. Therefore, we can set *PERIODICITY_HINT={52,12}*. Moreover, an event occurs every six weeks: either the growth stops, or there is a sudden fall in sales. Therefore, we can hint the forecast method even more by setting *PERIODICITY_HINT={52,12,6}*.


The following illustration shows the forecast when only the hint is considered and periodicity is not automatically detected (*AUTO_DETECT_PERIODICITY=0* and *PERIODICITY_HINT={52,12,6}*).



So far, this seems to be the best prediction: the sales are on a descending trend, so they can't be as good as they were during the previous year, but they still follow the same seasonality as the previous year.

It's important to note that even if you provide the model a hint, the *AUTO_DETECT_PERIODICITY* parameter still plays a role. However, the better the hint, the less impact the *AUTO_DETECT_PERIODICITY* parameter has on the forecast.

The following illustration shows the forecast when *AUTO_DETECT_PERIODICITY*=1 and *PERIODICITY_HINT*={12}.



This is quite different from what we got when we used *AUTO_DETECT_PERIODICITY*=0 and *PERIODICITY_HINT*={52,12,6}.

If we use *AUTO_DETECT_PERIODICITY*=1 and *PERIODICITY_HINT*={52,12,6}, notice that we get the same forecast as when we use *AUTO_DETECT_PERIODICITY*=0 and *PERIODICITY_HINT*={52,12,6}. This proves that many periodicity hints likely lead to significantly longer model training times but more accurate models.


FORECAST_METHOD

To control the choice of algorithm, use the *FORECAST_METHOD* parameter. By default, it's set to MIXED, which means that both ARIMA and ARTxp are used in the prediction, and both have equal weight.

Note: In SQL Server Enterprise Edition and SQL Server BI Edition, users can customize the blend by selecting the **MIXED** option and then setting a parameter (*PREDICTION_SMOOTHING*) that weights the results toward either the short-term prediction or the long term prediction.


The following illustration shows the result in DemoData for item D0003 if you use only the ARIMA method. The forecast has been generated in the same way as before:

- **Historical horizon start date** = 01 January 2012
- **Historical horizon end date** = 01 December 2012
- **Baseline forecast start date** = 02 December 2012
- **Forecast bucket** = Week
- **Forecast horizon** = 12



The following illustration shows the result in DemoData for item D0003 if you use only the ARTxp method. The forecast has been generated in the same way as before:

- **Historical horizon start date** = 01 January 2012
- **Historical horizon end date** = 01 December 2012
- **Baseline forecast start date** = 02 December 2012
- **Forecast bucket** = Week
- **Forecast horizon** = 12



Also notice that the prediction is for only three weeks, not 12. This is because Analysis Services stops the ARTxp prediction if it identifies that it becomes too noisy – in other words, the predicted values have peaks that are too high and falls that are too deep. The *INSTABILITY_SENSITIVITY* parameter controls the point at which ARTxp predictions are stopped.

So far, we have looked at parameters that control seasonality detection. The next two parameters control the instability in predictions.


INSTABILITY_SENSITIVITY

Because of the long-term instability of ARTxp, the Analysis Services development team implemented code that predicts when the forecast will grow unstable and stops returning values. This functionality was based on the assumption that no values are better than wildly out-of-range values. Customer feedback proved this assumption to be wrong, so this parameter was introduced to let users control when ARTxp stops the prediction.

This parameter controls the point at which prediction variance exceeds a certain threshold. After that, the ARTxp algorithm suppresses predictions. This implies that the parameter is valid only when *FORECAST_METHOD* is set to ARTxp. The default value of 1 is the strongest cutoff for instability. If the parameter is set to 0, instability detection is turned off.


For example, we don't want the ARTxp prediction to stop for D0003 and therefore sets *INSTABILITY_SENSITIVITY*=0.

The following illustration shows the forecast that is created.



Notice that the prediction continues all the way to 12 weeks ahead.

The following illustration shows the forecast when we set the prediction threshold to 0.1 (*INSTABILITY_SENSITIVITY*=0.1).



This parameter can be modified only in SQL Server Enterprise Edition. In SQL Server Standard Editions, Analysis Services uses only the default value of 1. Therefore, Analysis Services monitors the normalized standard deviation for each prediction, and as soon as this value exceeds the threshold for any prediction, the time series algorithm returns NULL and stops the prediction process.


PREDICTION_SMOOTHING

This parameter is available only in SQL Server Enterprise Edition. It controls how *FORECAST_METHOD=MIXED* works, by letting users give more weight to one of the algorithms. As has already been mentioned, when *FORECAST_METHOD=MIXED*, the models are trained with both algorithms, and then the results are blended. During this process, both algorithms have equal weight. Therefore, this parameter has an effect only if *FORECAST_METHOD=MIXED*, and it is responsible for controlling instability.

PREDICTION_SMOOTHING takes a value between 0 and 1. A value of 0 means that only ARTxp is used, and a value of 1 means that only ARIMA is used. If any other value between 0 and 1 is used, an exponential decreasing function of steps is created, and is used to fade out ARTxp and cross-fade in ARIMA.


The illustrations that follow are taken from the Microsoft Time Series Algorithm Technical Reference and show this exponential function.

The following illustration shows the exponential function when *PREDICTION_SMOOTHING=0.5*.



For the first forecast bucket, ARIMA and ARTxp are weighted the same (0.5 each); but as the number of forecast buckets increases, ARIMA get more weight, trending to 1 at around 46 to 49 forecast buckets. Meanwhile, ARTxp loses weight, trending to 0 at around 46 to 49 forecast buckets.

The following illustration shows the exponential function when *PREDICTION_SMOOTHING*=0.2.



For the first time buckets, ARTxp has the greatest weight ($1-0.2=0.8$), but this weight decreases, trending toward 0.1 as we reach between 46 and 49 buckets. Meanwhile, ARIMA has a small weight in the forecast of the first buckets (0.2) but gets more weight, trending to 0.9 when we are above 46 buckets.


In DemoData, let's look again at item D0003 to consider an example in which the methods are blended.

Generate a forecast with the following parameters:

- **Historical horizon start date** = 01 January 2012
- **Historical horizon end date** = 01 December 2012
- **Baseline forecast start date** = 02 December 2012
- **Forecast bucket** = Week
- **Forecast horizon** = 3

Because the forecast horizon is less than 5, we are using ARTxp as the only forecast method.

The following illustration shows the result.




This is quite different from the long-term prediction we got when using only ARIMA (see the illustrations for the *PERIODICITY_HINT* parameter). The prediction is not smoothed out.

Now increase the forecast time fence to 6 buckets, and start blending the algorithms.

The following illustration shows that when *FORECAST_METHOD*=MIXED and *PREDICTION_SMOOTHING*=0, we get the same result as before for the first three weeks.




The following illustration shows that using *FORECAST_METHOD*=MIXED and *PREDICTION_SMOOTHING*=1 is basically like using ARIMA, so the prediction is smoothed out:



It seems like the prediction is smoothed too much, especially for the next two weeks, so let's give both algorithms equal weight (*FORECAST_METHOD=MIXED* and *PREDICTION_SMOOTHING=0.5*).

The following illustration shows the result of this change.



As you can see, the forecast has basically become a blend of the two previous forecasts: it's flattened out because of ARIMA, but it keeps the "shape" of ARTxp.

MINIMUM_SERIES_VALUE

This parameter specifies the minimum value that can be predicted. It is used together with *MAXIMUM_SERIES_VALUE* to constrain the predictions to some expected range. For example, it can be used to specify that the predicted sales quantity should never be a negative number.

This parameter is available only in SQL Server Enterprise Edition. However, AX 2012 R3 lets you use this parameter with any SQL Server edition. Support for this parameter was added because of the known Excel issue with weighted allocation when cells that have a value of 0 are modified, as described earlier in this white paper.

Note that we apply this parameter even when we generate a forecast by using the **Copy over historical demand** forecast strategy.

MAXIMUM_SERIES_VALUE

This parameter specifies the maximum value that can be predicted. It is used together with *MINIMUM_SERIES_VALUE* to constrain the predictions to some expected range. For example, it can be used to specify that the predicted sales

quantity for any day should never exceed the number of products in inventory.

This parameter is available only in SQL Server Enterprise Edition. However, AX 2012 R3 lets you use this parameter with any SQL Server edition.

Note that we apply this parameter even when we generate a forecast by using the **Copy over historical demand** forecast strategy.

MISSING_VALUE_SUBSTITUTION

It's very common to have gaps in the time series data, such as a bucket that has no historical data or a month that has no sales for a specific product.

Gaps inside a time series are allowed in Analysis Services.

The following illustration shows an example of a time series that is allowed even though it is missing the whole period between February 2013 and December 2013.

Read only historical and forecast quantity						
Row Labels	October 2012	November 2012	December 2012	January 2013	February 2013	December 2013
D_Aloc	1089	1701	647	117	73	1
D0001	378	390	158	39	23	1

However, by default, gaps caused by alignment between time series data are not allowed in time series in Analysis Services. Therefore, this parameter specifies how these gaps in historical data should be filled.

The following illustration shows an example of a misalignment in time series data.

Read only historical and forecast quantity								
Row Labels	August 2012	September 2012	October 2012	November 2012	December 2012	January 2013	February 2013	December 2013
D_Aloc	995	1028	1089	1701	647	117	73	1
D0001	344	355	378	390	158	39	23	1
D0002				574	202			
D0003	344	355	380	387	145	42	24	
D0004	307	318	331	350	142	36	26	

In this example, item D0002 is missing time series information both from August 2012 to December 2012 and from January 2013 onward. However, it has to be forecasted together with items D0001, D0003, and D0004, which have data points in all the buckets.

Because such gaps are hard for users to catch, we decided to improve the usability of demand forecasting in AX 2012 R3 by filling these alignment gaps with 0 values.

Before we explain this parameter in greater details, note that the process of "filling" missing values is more complex than you might expect. There can be physical filling and mining model filling of data gaps. Depending on what type of filling you choose to do, the performance of demand forecast generation and the accuracy of the forecast can vary quite significantly.


Physical filling means that the data gaps are filled with 0 values in the database. However, no Microsoft Dynamics AX transactional data is affected. The values are saved only in the tables used for demand forecasting calculation. This type of substitution is done by AX 2012 R3. To use it, select the **Do not allow gaps in historical data** check box in the **Demand forecast parameters** form.

Mining model filling means that the gaps are never really filled and stored in the database; instead, during training, the mining model uses the value provided by the *MISSING_VALUE_SUBSTITUTION* parameter. This is the native way of dealing with data gaps in Analysis Services. Mining model filling gives better performance when a demand forecast is generated, and it can also offer better forecast accuracy. On the other hand, our tests have shown that it's not that easy to understand the results it produces.

To illustrate the difference between the two filling modes, generate a forecast with the following settings:


- Historical horizon start date = 01 January 2012
- Historical horizon end date = 30 June 2012
- Baseline forecast start date = 05 August 2012
- Forecast bucket = Week
- Forecast horizon = 12
- FORECAST_METHOD=ARIMA

The following illustration shows the forecast for item D0003 when the **Do not allow gaps in historical data** check box is selected.



Every gap in demand is replaced with 0, which suggests to the forecasting method that the demand should decrease. The seasonality is kept, but the 0 values have lowered the forecast.

The following illustration shows the forecast for item D0003 when the **Do not allow gaps in historical data** check box is cleared.



It looks quite different: the forecast is not lowered in any way.

Users can specify which type of filling to use by setting the **Do not allow gaps in historical data** check box in the **Demand forecasting parameters** form, under **Master planning > Setup > Demand forecasting**. If this check box is selected, physical filling of historical data gaps is performed.

The *MISSING_VALUE_SUBSTITUTION* parameter is used only when the **Do not allow gaps in historical data** check box is cleared, and it takes several values:

- **None** (the default value)
- **Mean** – Uses a moving average of time buckets and smooths out the history.
- **Numeric constant** – Uses a specified number to replace the missing values.
- **Previous** – Repeats the value from the previous time bucket.

In DemoData, make sure that the **Do not allow gaps in historical data** check box is cleared. We can then look at an example of the native Analysis Services method for missing value substitution for item D0003.

Generate a forecast with the following parameters, also making sure that items D0001, D0002, D0003, and D0004 are part of the forecast:


- **Historical horizon start date** = 01 January 2012
- **Historical horizon end date** = 31 December 2012
- **Baseline forecast start date** = 01 January 2014
- **Forecast bucket** = Month
- **Forecast horizon** = 12

We are creating the forecast for the year 2014, based on the forecast for year 2012, with year 2013 missing in between.

MISSING_VALUE_SUBSTITUTION=None should result in the following error: "Time stamps not synchronized starting with series, of the mining model, Forecast model – [xxxxxxx]. All time series must end at the same time mark and cannot have arbitrarily missing data points. Setting the *MISSING_VALUE_SUBSTITUTION* parameter to Previous or to a numeric constant will automatically patch missing data points where possible."

To improve the user experience, we fill in the alignment gaps with 0 values, so users will never receive this error message in AX 2012 R3. We are doing this because, in this case, you can't use any value other than 0 to fill in the data gap. Any other value for *MISSING_VALUE_SUBSTITUTION*, such as Mean or Previous, won't make sense, because data cannot be generated out of nothing.

The following illustration shows the result when *MISSING_VALUE_SUBSTITUTION*=Mean, *PERIODICITY_HINT*={12}, and *FORECAST_METHOD*=ARIMA.




The 2014 forecast for item D0003, based on year 2012, is smoothed out to become an average of the sales in 2012. This

happens mostly because we are using ARIMA in this example. If we replace ARIMA with ARTxp, the result will be so unstable that Analysis Services suspends the prediction and no results are calculated for D0003.

There isn't much difference between using Mean, Previous, or a numeric constant as the value for *MISSING_VALUE_SUBSTITUTION*. The gap is there, so not much can be done. What is interesting is how Analysis Services distributes the predicted values, based on the gaps in the history. To see this, generate a forecast with the following settings:

- **Historical horizon start date** = 01 January 2012
- **Historical horizon end date** = 30 June 2012
- **Baseline forecast start date** = 05 August 2012
- **Forecast bucket** = Week
- **Forecast horizon** = 12
- **FORECAST_METHOD**=ARIMA

The following illustration shows the historical transactions and forecast for item D0003.




Notice that there are gaps in the historical demand for D0003: the blue bars don't come immediately one after the other; by contrast, there are no gaps in the forecast: the red bars are continuous. At the same time, as the trend lines show, the forecast more or less replicates the historical values.

As a general rule, the forecast never has gaps. If Analysis Services doesn't return a forecast for a bucket, the bucket is added anyway, and the *MINIMUM_SERIES_VALUE* parameter is used to fill it in.

As a down side, inventory levels grow, because demand that was distributed over 16 weeks is now distributed over 4 weeks. To mitigate this, you can use physical filling of missing values. Select the **Do not allow gaps in historical data** check box, and then generate a forecast as before:

- **Historical horizon start date** = 01 January 2012
- **Historical horizon end date** = 30 June 2012
- **Baseline forecast start date** = 05 August 2012
- **Forecast bucket** = Week
- **Forecast horizon** = 12
- **FORECAST_METHOD**=ARIMA

The following illustration shows the historical transactions and forecast for item D0003.



In this case, the forecast demand is spread, but also notice that the demand forecast is leveled out and the accuracy seems to decrease.

COMPLEXITY_PENALTY and MINIMUM_SUPPORT

These two parameters apply to the ARTxp method only, so they have effect only when *FORECAST_METHOD*=MIXED or *FORECAST_METHOD*=ARTXP.

COMPLEXITY_PENALTY is available only in Analysis Services Enterprise Edition.

It's hard to grasp these two parameters without understanding the Microsoft Decision Tree algorithm. At the same time, it's hard to understand the Microsoft Decision Tree algorithm without having a rather deep knowledge of algorithms, data structures, and mining models. For details about the algorithm, see [Autoregressive Tree Models for Time-Series Analysis](#).

For more information about these two parameters, see the Internet resources in the [References](#) section.

Tips for setting up demand forecasting algorithm parameters

- The more history, the better. For the best accuracy, make sure that you have at least 40 data points.
Note: If the object you are forecasting changes quickly and becomes obsolete quickly, not many data points are needed.
- If the forecast horizon is more than five forecast buckets, do not use only the ARTxp forecasting method, because it tends to become noisy. But ARTxp is recommended by the Analysis Services team for predictions of less than five buckets.
- The *PERIODICITY_HINT* parameter has a huge influence on the forecast. It might be the most important one.
- If your data has periodicities, set them to be as granular as possible.
For example, for sales, set *PERIODICITY_HINT*={12,3,1} if the sales vary yearly, but there are events that happen every quarter in a year and every month in a quarter.
If you use such granular settings, model training performance will not be as good, but forecast accuracy will be better.
- After you change the forecast bucket, remember to change the periodicity hint, because it might not make sense anymore.
- If you get forecasts that are very noisy – that is, they have lots of peaks and sudden falls – ARTxp is having too much influence on them. Smooth them out by using ARIMA.
- For better performance, leave the **Do not allow gaps in historical data** check box cleared.
- Analysis Services creates one mining model for each item allocation key. This means that, for better forecast accuracy, products that are grouped together in one item allocation key are forecasted together (correlations are established among the products). This is why it is very important to group products that have similar demand patterns (similar base values, trends, and seasonality) in the same item allocation key.
Note: Unfortunately, this setup tip applies only when you use SQL Server Enterprise Edition or BI Edition. SQL Server

Standard Edition allows for a single model only.

References

- Microsoft Time Series Algorithm Technical Reference – <http://technet.microsoft.com/en-us/library/bb677216.aspx>
- <http://www.forecastingmethods.net/Demand-Forecasting-Methods.html>
- <http://sqlmag.com/sql-server-analysis-services/understanding-time-series-forecasting-concepts>
- TechNet webcast – <http://msevents.microsoft.com/cui/WebCastEventDetails.aspx?EventID=1032359362&EventCategory=4&culture=en-US&CountryCode=US>
- *Data Mining with Microsoft SQL Server 2008*, by Jamie MacLennan, ZhaoHui Tang, and Bogdan Crivat, John Wiley & Sons, 2009
- <http://social.msdn.microsoft.com/Forums/sqlserver/en-US/f9e85016-6d01-4215-9942-5c6c81f88f4a/complexitypenalty-parameter-and-forward-pruning-in-decision-trees-sql-server-2008?forum=sqldatamining>

Performance improvement tips and tricks

Through the development process, we have seen that, because of the data-intensive nature of forecasting, users might encounter certain performance challenges. The following sections provide guidelines about potential usage patterns, and explain how to identify parameters and contributing factors that affect the performance and run times. We hope these guidelines will help you make conscious decisions while being aware of the potential impact of these changes.

Important! Note that this document does not address any specific performance issue. Instead, it provides multiple recommendations and pointers about what factors can adversely affect the performance and run time of the feature set.

A brief technical look at the Demand forecasting implementation

The core functionality for demand forecasting can be summarized as consisting of three essential operations:

1. **Forecast generation** – This is the most complex operation, and we can summarize the steps as follows:
 1. Collecting historical data through X++.
 2. Processing and preparing the data so that it can be passed to the Analysis Services service. This applies only if statistical forecasting is used. It does not apply if historical data is copied to a new time horizon.
 3. Training the Analysis Services data mining structure and models. This applies only if statistical forecasting is used. It does not apply if historical data is copied to a new time horizon.
 4. Reading back the forecast from the Analysis Services prediction engine. This applies only if statistical forecasting is used.
 5. Processing the cubes to surface the data in Excel pivot tables.
2. **Forecast import** – After users have finished modifying the forecast, they submit it back to the cube and then import it back into Microsoft Dynamics AX.
3. **Forecast accuracy calculation** – This functionality compares the actual demand in a given period against the predicted demand, and then measures the accuracy of the predictions.

All the surrounding functionality enables, controls, or monitors these three operations. Forecast generation is by far the most complex and data-heavy operation, and is the primary focus of the following sections.

Hardware recommendations

Because demand forecasting is a data-intensive operation, many hardware components come into play and might affect its performance. A quick summary follows:

- **Disk IO** – Because of the number of data read/write operations that take place, good I/O speed is important.
- **CPU** – Although multi-threading is currently not supported by Demand forecasting on the X++ side, it is important that the processing power be adequate for computational operations and possible performance gains on the SQL and Analysis Services operations where parallelization is used.
- **Memory** – It was observed that the most direct effect on performance is through the memory constraints. When total memory consumption approaches around 95 percent of the total available memory, many processes start

swapping data between the memory and the disk so that they can continue running, which results in much higher run times. The following processes have been observed to be critically affected when memory is constrained:

- **Analysis Services data mining training** – Low memory might result in exponential run times.
- **Cube processing** – Low memory might result in exponential run times.
- **Excel pivot tables, drill-through** – When dimensions are added to the pivot table for drill-through, lower memory resulted in exponential run times.

Important! If you start experiencing a sudden slowdown of any processes that are running as a part of demand forecasting, make sure that memory consumption is not approaching the memory limit of the system. If it is, consider adding memory to the system.

System setup

We recommend that you set up Microsoft Dynamics AX as a three-box environment with dedicated Application Object Server (AOS) and SQL Server machines, and with external client connections.

For demand forecasting, the same recommendation applies but with an additional factor. By running Analysis Services on a dedicated machine you might improve performance, which can improve run times. This is because a dedicated environment guarantees that most of the resources are allocated to Analysis Services itself, so it doesn't have to compete for resources, especially with the SQL Server instance or with any other OS process.

SQL Server and Analysis Services editions

The Demand forecasting feature works with all the non-Express editions of SQL Server from SQL Server 2008 SP1 onward. Therefore, in terms of the functionality, there are no limitations on the edition that you can use.

However, it is important to know that SQL Server Standard Edition does not support multiple mining models in Analysis Services, so you cannot train mining modes or obtain predictions in parallel. Therefore, from the performance point of view, we recommend that you upgrade to higher editions of SQL Server if you want to process larger datasets. The following table describes the SQL Server editions that support parallel mining models.

SQL Server edition	This edition supports multiple mining models
SQL Server 2008 SP1 Standard Edition	No
SQL Server 2008 SP1 Enterprise Edition	Yes
SQL Server 2008 SP1 Data-warehouse Edition	Yes
SQL Server 2012 Standard Edition	No
SQL Server 2012 Enterprise Edition	Yes
SQL Server 2012 BI Edition	Yes
SQL Server 2014 Standard Edition	No
SQL Server 2014 Enterprise Edition	Yes
SQL Server 2014 BI Edition	Yes

A deeper look at the steps for statistical forecast generation

As we mentioned in the introduction of this section, statistical forecast generation is the most complex operation of the feature set, with many individual sub-tasks taking place. We will explain some of the major tasks in more detail in this section to provide more insight and set the context for some of our suggestions.

In a typical breakdown of statistical forecast generation, many tasks occur.

Tip: All of these individual tasks and details are logged under **Master Planning > Periodic > Demand forecasting > Session log**.

Depending on the size of the dataset, most tasks are not time consuming. However, in larger datasets, we can identify two main operations as the heaviest operations to run. These are the training of the mining models and import of the forecast. These are described in the following table.

Stage	Description
Prepare the data for the data mining models.*	<p>During this stage, the data that should be passed for mining model training (TSQL query) takes place. We can summarize the events as follows:</p> <ol style="list-style-type: none"> 1. Historical data is collected and expanded based on the forecasting dimensions, which create the identifying granularity keys (for example, [DataAreaID][Item allocation key][ItemID][Color]). Important! The individual granularity keys generated here go to Analysis Services as individual time series that can be used for forecast generation. The greater the number of unique granularity keys, the more time series, and therefore, the more processing time. 2. The data is aggregated based on the date dimension, matching the bucket type (month, week, day). Important! Manual missing value substitution also takes place here if the parameter is set by the user. This essentially adds zeroes (0) as values for each possible combination of forecasting dimensions (granularity key) where no data previously exists, which could result in an explosion of data size. 3. Date alignment: For all unique combinations of granularity keys that contain data, a 0 is added in the first bucket of the historical horizon and in the last bucket before the forecast start date. This is done to align the forecast, because Analysis Services starts the forecast immediately where the historical data ends; therefore, the forecast might end up with fewer points than required. For example, if the historical data ends three months before the forecast start date, and we are forecasting for six months in the future, we might only get three months of forecast, because the algorithm returns six data points, with three in the months prior to the forecast start date.
Create and process the data mining structure.	<p>This step primarily prepares the Analysis Services mining structure and models. In this step, the following events take place:</p> <ol style="list-style-type: none"> 1. For the SQL editions that support multiple mining models, for a single mining structure, a different mining model is generated for each item allocation key. 2. For each granularity key, a unique time series will be generated. 3. The actual training and processing of the mining structure and mining models take place.
Import the demand forecast from the data mining models.	<p>This step can be summarized as follows:</p> <ol style="list-style-type: none"> 1. MDX queries are used to get the next N predictions. Tip: The actual forecast predictions take place in this step, which might result in longer run times. 2. The predictions are imported back into Microsoft Dynamics AX tables (OLAP to OLEDB), and number range normalizations also take place.

* We have included this stage in the table to point out that it directly affects the run times for the next steps in the chart.

Performance improvement tips for forecast generation

Now that the important steps are clear, we can have a look at some tips.

Note: These tips are mostly related to forecast generation, but many also have indirect impact on forecast import, and on accuracy generation, which shares some steps with forecast generation.

- **Use fewer forecasting dimensions over sparse data** – If many dimensions are enabled for drill-through (such as Stock keeping unit, Customer group, Customer account, and Country/region), more unique granularity combinations are generated, and Analysis Services must generate, train, and process many more time series. The

performance impact of this is obvious, and will result in a proportional and sometimes even exponential increase in run time for the previously mentioned tasks and others.

Tip: If you have large data volume and many unique forecasting dimension combinations in your transactional data, try to avoid forecasting on unnecessary dimensions, or expect higher run times.

- **Use only the relevant part of the historical data** – It is self-explanatory that the more data that is processed, the longer the run times.

Tip: Try to restrict your historical data to the necessary portion, which is usually close to the current date. For example, don't base your forecast on the past 10 years of history if only the last year is relevant. This might also result in a more accurate forecast.

- **Limit the forecast horizon** – The longer the forecast horizon, the more prediction points need to be generated for individual unique granularity keys for the forecast, which result in extended run time.

Tip: Try to keep your forecast horizon limited to the time frame that fits your business needs to avoid unnecessary processing.

- **Use multiple item allocation keys:** For the SQL Server editions that support multiple mining models, the mining model creation, training, and processing step can be parallelized by using multiple mining models. If you use only a single item allocation key with all the items that need to be forecasted on, you might miss out on the benefits of parallelization. We observed that a single mining model was much slower than multiple mining models in our internal test environments.

Tip: Try to use mining model parallelization by creating multiple item allocation keys if applicable.

Important! For the items under the same item allocation key, Analysis Services seeks out correlations while creating the time series. Therefore, it is important for forecast accuracy to assign items that have a similar trend and seasonality in the same item allocation key. Try to ensure this for the best results.

- **Manual vs. Analysis Services default missing value substitution** – When you have gaps in your historical data, there are two methods for filling the missing values:

- **Default Analysis Services missing value substitution** – Uses the missing value substitution parameter for Analysis Services mining models.

- **Manual replacement** – Generates a 0 value for each unique granularity key where there is no data. This method can sometimes result in a more accurate and stable forecast.

Tip: We recommend the default method in most cases. Manual replacement of missing values should be used only when deemed functionally necessary, because it might result in performance issues, and cause lots of redundant data to be generated and passed to Analysis Services. Use with caution.

- **Running in batch mode** – Because the SysOperation framework is used, the data heavy operations always run in batch mode behind the scenes. However, if you use the run-in-batch method, your client is not locked for the duration of the operations.

Tip: Forecast generation, import, and accuracy calculation always runs in a batch, even if you do not select **run in batch**.

- **Copy over forecasting vs. statistical forecasting and blending the two methods** – Copy-over is the forecasting method that simply copies the historical forecast to the new forecast horizon. Because no forecast algorithm or heavy processing is involved, it is usually much faster at run time.

Although it is possible to run forecasting by using either one of the methods for all item allocation keys, it is also possible to blend these methods together. This is done by setting the forecast generation method for individual item allocation keys in the **Demand forecasting parameters** form (see Help for more information).

Tip: In many cases, the copy-over method is faster. See whether you can use it for some or all of your item allocation keys.

Demand forecast performance cheat sheet

The following list summarizes some of the tips mentioned in this document. You can use this as a checklist during setup or when you want to improve performance.

- Always make sure that you have enough memory/RAM.
- A dedicated Analysis Services server is a good idea.
- Choose the SQL Server edition that fits your business needs.
- Enable only the forecasting dimensions that you need.
- Use only the portion of the historical data that you need and trust.
- Limit the forecast horizon to what is necessary.

- Try to use mining model parallelization by correctly distributing your items to item allocation keys.
- Assign only items with similar characteristics to the same item allocation key.
- Take careful consideration before using the manual method to insert the missing value substitution for gaps in data.
- The forecast runs in batch mode (IL) even if you don't select **run as batch**, but you can use the batch mode to avoid locking the client.
- Consider using the copy-over forecast functionality where applicable, either for all item allocation keys or for at least for some of them.

Send feedback.

Microsoft Dynamics is a line of integrated, adaptable business management solutions that enables you and your people to make business decisions with greater confidence. Microsoft Dynamics works like and with familiar Microsoft software, automating and streamlining financial, customer relationship, and supply chain processes in a way that helps you drive business success.

United States and Canada toll free: (888) 477-7989

Worldwide: (1) (701) 281-6500

www.microsoft.com/dynamics

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, this document should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation. Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2014 Microsoft. All rights reserved. Microsoft, Microsoft Dynamics, and the Microsoft Dynamics logo are trademarks of the Microsoft group of companies.