WHITE PAPER

NICE WFM 7.0:
Forecasting with Artificial Intelligence

CONTACTS

Global International HQ, Israel,
T +972 9 775 3777, F +972 9 743 4282

Americas, North America
T +1 551-256-5000, F +1 551-259-5252

EMEA, Europe & Middle East,
T +44 0 1489 771 200, F +44 0 1489 771 665

Asia Pacific, Singapore Office
T + 65 6222 5123, F +65 6222 5459

With NICE (Nasdaq: NICE), it's never been easier for organizations of all sizes around the globe to create extraordinary customer experiences while meeting key business metrics. Featuring the world's #1 cloud native customer experience platform, CXone, NICE is a worldwide leader in AI-powered self-service and agent-assisted CX software for the contact center – and beyond. Over 25,000 organizations in more than 150 countries, including over 85 of the Fortune 100 companies, partner with NICE to transform - and elevate - every customer interaction.

www.nice.com

TABLE OF CONTENTS

INTRODUCTION

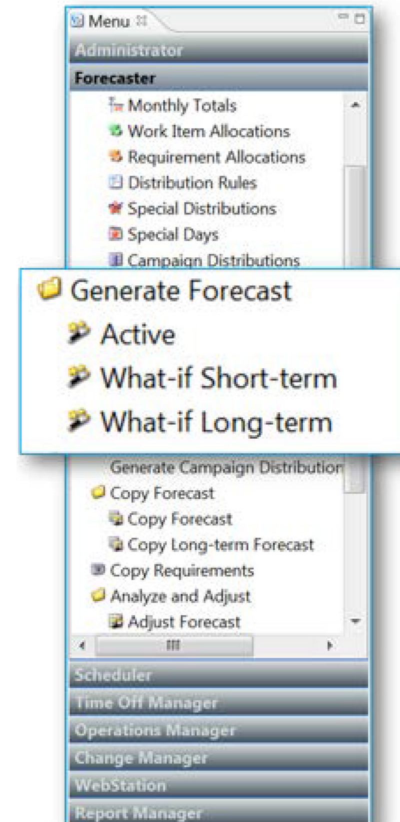

We are living in a golden age of predictive analytics, and forecasting and predictive technologies are revolutionizing business well beyond the contact center. These forecasting capabilities are supported by ongoing developments in machine learning and artificial intelligence, developments that are enabling increasingly accurate projections.

The most robust contact center forecasting tools on the market have far-reaching impact. Advanced statistical methods can help users realize consistent customer service, improve retention and lower costs across the board through market-leading capabilities.

Although the concept is simple, the execution is anything but. A precise forecast requires an understanding of traffic during normal, day-to-day operations, but it must also take into account special considerations, such as spikes due to back-to-school shopping season, a new marketing promotion or a change in operating hours. A precise forecast is the foundation on which scheduling is built.

NICE WFM 7.0's Forecaster unlocks a high level of transparency into interaction history, allowing you to centrally forecast, schedule and manage contacts between multiple locations and ensure that site- and enterpriselevel objectives are met. With more than two thousand customers and two million users depending on its unparalleled ability to fine-tune the most precise forecasts, Forecaster allows you to plan and respond to the peaks and valleys of customer history through automatic collection of key historical data from all types of contact sources:

- Automatic call distributors (ACDs)
- Outbound dialers
- Multi-channel routing platforms
- Back-office employee desktops

NICE WFM 7.0 offers three types of forecasts to ensure that your call center's unique needs are met:

- active, which is used for allocations and scheduling;
- what-if short-term, which enables you to model scenarios in 15- and 30-minute increments to review the effects of proposed changes to the plan; and
- what-if long-term, which enables future strategic planning and budgeting through scenarios in day, week or month increments.

Forecasting 101

It's an often-repeated truth that creating an accurate forecast is one part science, one part art. And science is in fact a key component of generating a forecast in which you can trust. When forecasting demand in a contact center, back office, branch or retail organization, we can draw upon a number of techniques from mathematics and statistics.

The challenge, however, lies in having the expertise needed to understand which forecasting method will generate the most accurate results in any given situation. A method that works effectively for one work stream may fail to perform well for another, or it may work sometimes for a work stream but not all of the time. Other considerations include operational context, the stage of a work stream's life cycle, the availability of historical data, the relevance of history to the future and the time horizon being forecasted.

The sheer number of forecasting techniques can seem overwhelming. There are, however, three basic types: time series, causal and qualitative.[1]

- Time-series methods focus on pattern recognition and pattern changes and rely solely on historical data. They perform well when trends or cycles are evident and stable.
- Causal methods rely on established cause-and-effect relationships between the data to be forecasted and other factors that influence the data.
- Qualitative methods are used when data is scarce. These methods use judgment to turn qualitative insights into quantitative estimates.

Most workforce management solutions on the market utilize a single forecasting method, but as of NICE WFM 7.0, the WFM forecaster has several forecasting methods available. Most of them employ time-series techniques, and one leverages a causal method using relationships it uncovers in historical data. No NICE WFM 7.0 forecasting models use the qualitative method due to the abundance of data available to the modern contact center.

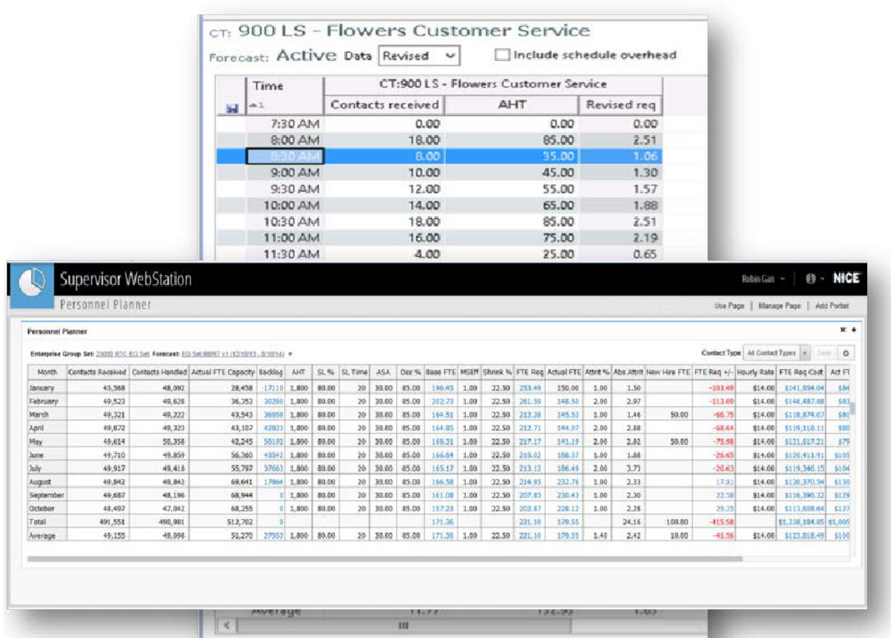

Figure 2: Sample adjust forecast screens

1. https://hbr.org/1971/07/how-to-choose-the-right-forecasting-technique
* Additional Online Resource: Online, Open-Access Textbooks - Forecasting: principles and practice: https://www.otexts.org/fpp/

In the sections that follow, we will briefly describe each method, including its uses, benefits and constraints. The descriptions are not intended to be exhaustive; to truly gain a complete understanding of each method, the user would need a background or education in statistical analysis. At the end of this paper, we provide a summary chart to help you consider the applicability of each method to each of your contact types.

Forecasting in NICE WFM 7.0

Forecasting requires that past historical data be collected and analyzed. Many factors, including (but not limited to) competitive activity, billing processes, weather, holidays, customer demographics, marketing activities and the inherent nature of the business, can affect the reliability and accuracy of a forecast and how it would typically be distributed across a day, week, or month.

NICE WFM 6.x and earlier versions include a weighted moving average model, and NICE WFM 7.0 introduces several new forecast models.

These new forecast models include:

- Box-Jenkins ARIMA
- Exponential Smoothing
- Multilinear Seasonal Regression
- Best Pick

These new models may be used with most of the existing automatic forecast adjustment factors, such as including extreme values, week of month seasonal cycles and week of year seasonal cycles. “Week Weights” and “Rate of change” are not used by the new forecasting models in NICE WFM 7.0.

Figure 3: CT forecast defaults example

How WFM 7.0 Models Use Optional Parameters

The NICE WFM 7.0 Forecaster offers nearly two dozen options for parameters, and the user can configure each parameter by contact type. Additional considerations are noted below for some of the options.

	Weighted Moving Average	Box-Jenkins ARIMA	Exponential Smoothing	Multilinear Seasonal Regression
Week Weights	Yes	No	No[1]	No
Maximum Weeks of History Used	13	156	156	156
Forecast Unit	Interval	Day of Week or Week[2]	Day of Week or Week[2]	Day of Week or Week[2]
Intraday and Intraweek Distribution Method	Not applicable	See footnote[3]	See footnote[3]	See footnote[3]
Include Extreme Values	Yes[4]	Yes[4]	Yes[4]	Yes[4]
Include Week of Month Seasonal Cycle	Yes	Yes[5]	Yes[5]	Yes[5]
Include Week of Year Seasonal Cycle[7]	Yes	Yes[5]	Yes[5]	Yes[5]
Contacts Abandoned to Include in Forecast Generations	Yes	Yes	Yes	Yes
Adjustments to Contacts	Yes	Yes	Yes	Yes
Adjustments to AHT	Yes	Yes	Yes	Yes
Rate of Change	Yes	No[6]	No[6]	No[6]
Special Days	Yes	Yes	Yes	Yes
Special Distributions	Yes	Yes	Yes	Yes
What-if Short term	Yes	Yes	Yes	Yes
What-if Long term	Yes	Yes	Yes	Yes
Active Forecast	Yes	Yes	Yes	Yes
Used for Contact Value Forecast	Yes	Yes	Yes	Yes
Used for AHT Forecast	Yes	Yes	Yes	Yes

Footnotes

1. The model determines the weights.

2. When the queue history is less than 6 to 8 weeks old (depending on the day of week the forecast is generated and the start day of week for the entity), forecasting by Day of Week is done automatically regardless of the selected interval.

3. If the forecast unit is "Week", the week's forecast value is distributed to each day of week and interval within a day. If the forecast unit is "Day of Week," each day's forecast value is distributed to each interval within the day. Intraday and intraweek distributions are estimated from queue history using the "substitute value" method described in Footnote 4.

4. If the Weighted Moving Average model is selected, no outliers are excluded in the 13-week dataset. If it is not selected, outliers are excluded in the 13-week dataset for intervals that are greater than 2 standard deviations from the mean (with Bessel correction) of all values on the same day of week and same interval within the 13 weeks.

   If any of the Box-Jenkins ARIMA, Exponential Smoothing or Multilinear Seasonal Regression models are selected, no outliers are excluded in the dataset. If they are not selected, a 13-week moving window, not the entire dataset, is used to determine outliers. Values within the moving window that are consistently greater than 2 standard deviations from the mean (with Bessel correction) of all values on the same day of week and same interval within the 13-week moving window are replaced with a substitute value.

   "Consistently" means that the interval exceeds 2 standard deviations at least half of the time that the period is evaluated within the window as the window shifts to include new data.

   The "substitute value" is equivalent to a single order exponential smoothing forecast with alpha value of 0.23 of all prior values for the same day of week and interval. An alpha value of 0.23 weights the most recent week of history at 23 percent of the total, the thirteenth week at about 1 percent and all older weeks combined at about 3.3 percent (relatively insignificantly). Outliers are not included in the calculation of substitution values.

   All four models treat intervals missing from history or excluded by Special Day definitions as if they are outliers and exclude them from the history.

5. Although this option may be used, it is not recommended. The model attempts to consider trends and seasonal cycles. Enabling this option may overstate the effect of the trend or seasonal adjustment.

6. The influence of a rate of change is already considered within the model as it analyzes trends.

7. The "week of year seasonal cycle" will be based on the week of year for the Box-Jenkins ARIMA, Exponential Smoothing and Multilinear Seasonal Regression models. For the weighted moving average model, it is based on the month of year.

The Weighted Moving Average Model

NICE WFM 6.x and earlier versions provide an interval-specific Weighted Moving Average forecast algorithm, which is a time-series model. This forecasting model is also available in NICE WFM 7.0.

The Weighted Moving Average model is well-suited for stable historical data. The weightings allow a certain degree of control over the influence of unstable historical data. This model has demonstrated the best accuracy for determining intraday distributions in the near future.

Interval-specific

- Is generated based on each interval's unique history rather than a daily or weekly value allocated to each interval.
- Results in improved intraday curve fitting that requires no user intervention.

Intraday Distribution

daily total is the product of interval-specific forecast values

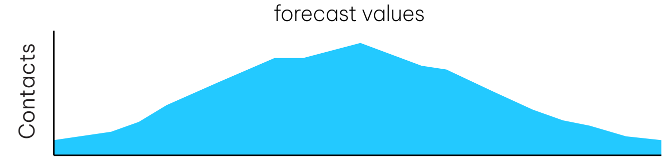

Weighted

- Uses weekly weights that are typically decreasing for older weeks, thereby reducing the influence of older data.
- Allows the user to influence the outcome by adjusting weights applied to the historical data.

Week	Demand	Weighting	Forecast
1	450	.5	
2	378	.3	
3	520	.2	
4			442.4
TOTAL:	1 348	1.00	

Moving average

- Creates a series of averages over time using a different subset of the full data set.
- Removes data that is older and less reliable by constantly moving the subset of historical data forward.

Intraweek Distribution

weekly total is the product of daily totals

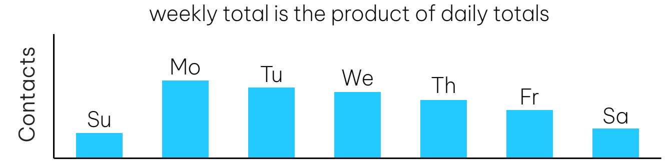

When to use:

- When the historical data has stationary patterns with or without trends or seasonality.
- When you can assume the past will continue to influence and represent the future.

Benefits:

- Easy to understand and implement.
- Simple to calculate.
- Provides stable forecasts.
- Smooths out short-term anomalies and can highlight long-term trends.

Constraints:

- May lag behind changes in trend if seasonal factors and other variables are not applied after the calculation.
- Complex relationships in data may be ignored.

The Box-Jenkins ARIMA Model

The Box-Jenkins ARIMA (Auto Regressive Integrated Moving Average) forecast model was introduced in 1970 by George Box and Gwilym Jenkins, who applied an iterative three-stage modeling approach to existing theory. It is a time-series method of forecasting that uses past values (the autoregressive model), past errors (the moving average model) or a combination of past values and past errors.

The Box-Jenkins method is well suited to handle complex time series forecasting situations in which the basic pattern is not readily apparent. It uses an iterative approach to identify a useful model from a general class of models.

The methodology does not make assumptions about any particular pattern within the historical data; instead, it iterates to identify patterns and incorporate them into the model. The chosen model—based on past values, past errors or a combination of both—is then checked against the actual historical data to see whether it accurately describes the history. If the model does not fit well, the process is automatically repeated until the most accurate model is found.

The basic approach to a Box-Jenkins ARIMA model is:

- Identification – Data is prepared and potential models are identified.
- Estimation and testing – Parameters in the potential model are estimated and analyzed until the best model is selected.
- Application – The selected model is used to produce the forecast.

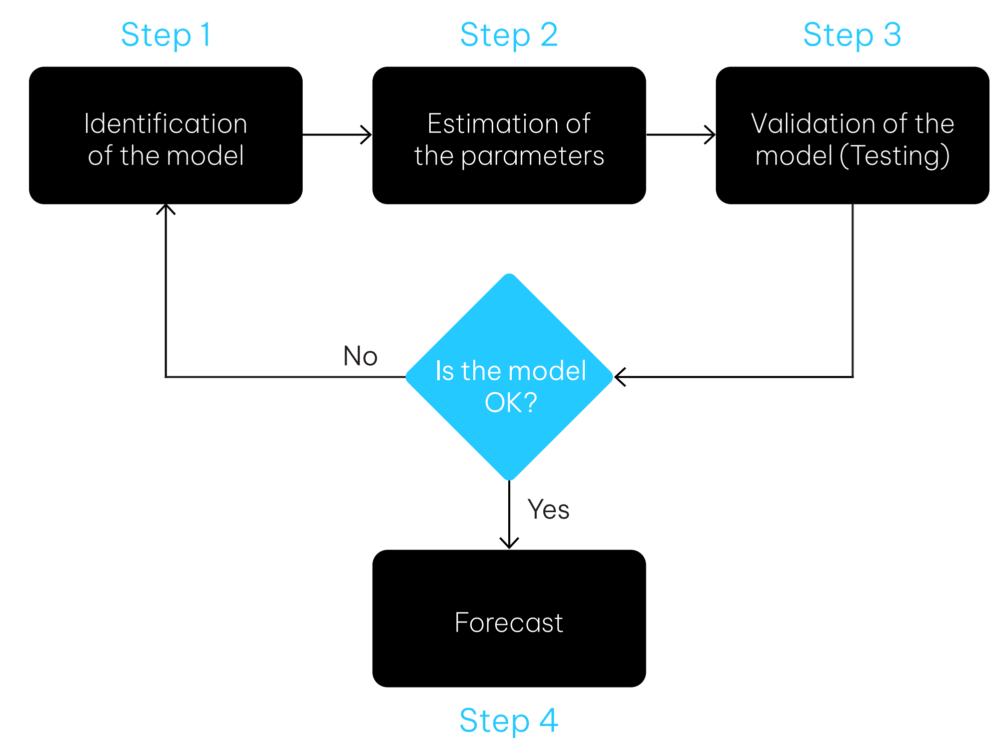

Since an ARIMA model is a type of time-series forecast, you can use it when you are able to assume that future patterns will be reasonably similar to those of the past. Because of this assumption, the model works best for short-term forecasting of 18 months or less and generally requires large historical data sets for improved accuracy.

A Box-Jenkins ARIMA model is similar to the Exponential Smoothing model in that both are adaptive and can model trends and seasonal patterns. In contrast to the Exponential Smoothing model, which is based on a structural view of level, trend and seasonality, ARIMA uses autocorrelations, or patterns in time.

Box-Jenkins performs better than Exponential Smoothing when the history is extended and more stable. Box-Jenkins is not well suited for wildly variable, more volatile or short history. In one study of 1,001 data sets, researchers found that Exponential Smoothing models outperformed Box-Jenkins models 55 percent of the time; Box-Jenkins was more accurate 45 percent of the time.[2] Given those results, it is worthwhile to consider switching between different techniques instead of assuming one method will always work all the time.

When to use:

- When there is stable data that has regular correlations.
- When the data patterns are very complicated, with a combination of trend, seasonal, cyclical and random fluctuationse.

Benefits:

- Is flexible and automatic in how it determines valid parameters.
- Can handle any data series, with or without seasonality.
- Works well for forecasts of 18 months or less.

Constraints:

- These models can be very sophisticated and may be difficult to explain.

2. S. Makridakis et al. [1984] The Forecasting Accuracy of Major Time Series Methods, Chichister: Wiley.

The Exponential Smoothing Model

Exponential Smoothing is a time-series model that provides a way to automatically calculate weights on all past data based on a determined smoothing factor. It is a more sophisticated approach than the Weighted Moving Average model.

Exponential Smoothing uses the average of historical data with exponentially decreasing multipliers. In its seasonal version, the Exponential Smoothing model uses three related equations for level, trend, and seasonality as the basis for a forecast. NICE WFM 7.0 includes single, double and triple Exponential Smoothing with nine different configurations: Multiplicative, Additive, and No Seasonality together with Linear, Damped, and No Trend.

The purpose of this method is to remove "high-frequency noise" (outliers) so that specific patterns can be determined in the historical data. This method is similar to the Weighted Moving Average in WFM 6.x and earlier in that both methods apply weights to the historical data. The model in WFM 7.0 is more advanced because the user does not have to supply the weights—the system determines the appropriate weights. The algorithm used in NICE WFM 7.0 is able to consider trend, seasonality and level (volume) simultaneously as it seeks the underlying pattern; the algorithm used in WFM 6.x applies a seasonality and trend factor after the base value is determined.

When to use:

- When historical data has stationary patterns with or without trends or seasonality.
- When you can assume the past will continue to influence and represent the future.
- When linear trend is present in the data.

Benefits:

- Provides good accuracy for short term forecasts.
- Gives recent data more weight.
- Adapts readily to changing patterns by automatically increasing or decreasing forecasts based on history.

Constraints:

- Can be less accurate for longterm forecasts because it can lag behind changes in trends.
- Bad data (anomalies) in most recent periods can affect the accuracy of the forecast.

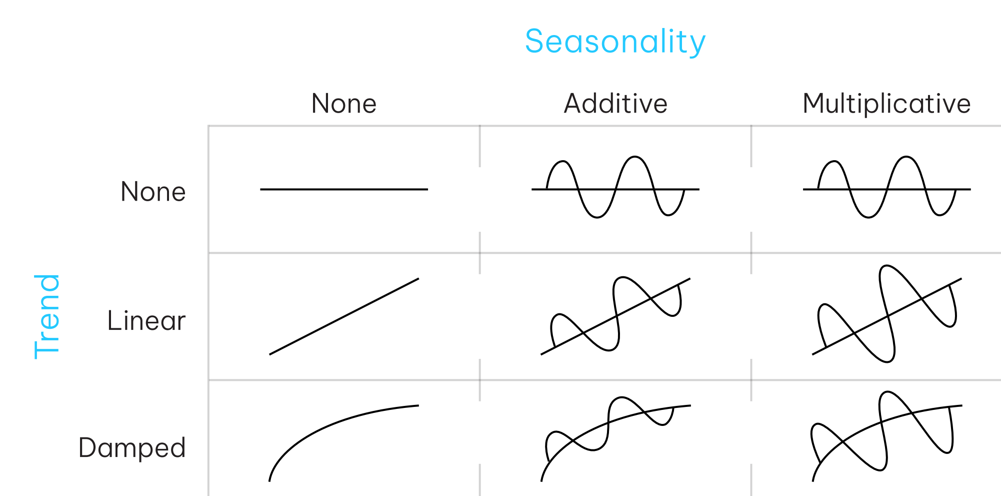

The Multilinear Seasonal Regression Model

Unlike the other time-series models used by the NICE WFM 7.0 Forecaster, Multilinear Seasonal Regression is an associative causal model used to model the trend and seasonality in a data set.

An associative causal model focuses on cause and effect. It assumes that the variable being forecasted is related to other known variables in the environment, and as a result, the forecasted variable is based on the associations to the known variables. If the known variables are actually "unknowable," then further assumptions are made to populate the variables with estimates in order to generate a forecast.

The NICE WFM 7.0 Multilinear Seasonal Regression model uses indicator variables in a multiple linear regression model to analyze seasonal effects. The indicator variables have two possible values, 0 or 1, depending upon whether the interval to which the variable applies is within the season that the variable indicates. Multilinear Seasonal Regression treats seasonal effects as additive rather than multiplicative.

When to use:

- When cyclical patterns exist but the rise and fall is not within a fixed period.
- When the future would be based on something other than time.

Benefits:

- Offers the ability to investigate the relationship between one or more variables to the criterion value of data.
- Is good for medium- to long-term forecasts.

Constraints:

- Incomplete data or limited historical data can lead to false correlations.
- Increased complexity makes it difficult to interpret or determine the best model.

The Best Pick Option

Unlike the other time-series models used If you are not sure which forecasting model to use, NICE WFM 7.0 provides an option for the system to select the best pick among the new forecast models: Box-Jenkins ARIMA, Exponential Smoothing and Multilinear Seasonal Regression.

If "Best Pick" is selected, the forecast generation will automatically generate and compare the output of all three models to identify the one that most closely matches historical data through a backcasting validation to find the minimum residual. If you use this method, you may avoid the iterative process of manually generating and reviewing the output of different models to find the best one to use.

When the forecaster selects the Best Pick option, NICE WFM 7.0 will evaluate each of the new models to:

- Determine the structure and parameters for each of the models.
- Compare the forecast accuracy of each model.
- Select the model that has the highest accuracy.

When to use:

- When the forecaster would prefer to have the NICE WFM system select the best forecast model.

Benefits:

- Reduces manual work for the forecaster.
- Allows automatic comparison of multiple forecast models in a single generation with the best model selected as a result of the generation.

Constraints:

- None.

Comparison Between Forecasting Models in NICE WFM 7.0

	Time Series			Causal
	Weighted Moving Average	Box-Jenkins ARIMA	Exponential Smoothing	Multilinear Seasonal Regression
Description	Each point of a moving average is the weighted mean of a number of consecutive points in the series, where the number of data points is chosen so that the effects of seasonality or irregularity or both are eliminated. This model smooths out short-term anomalies and can highlight long-term trends.	The time series is fitted with a mathematical model that is optimally derived through iterative estimating and testing of specific parameters to identify patterns.	This model is similar to the weighted moving average, except seasonal and trend effects are also computed.	The effect of seasonality is derived from history to establish a cause-and-effect relationship.
Accuracy				
Short Term (0-6 months)	★★★★	★★★★★	★★★★★	★★★★
Medium Term (6-18 months)	★★★	★★★★	★★★★	★★★★★
Long Term (18+ months)	★★★	★★★	★★★	★★★
Identification of Turning Points*	Poor	Good	Fair	Very good
Typical Applications	Low- to high-volume queues with stable historical data.	Medium- to high-volume queues with no readily apparent patterns.	Low- to high-volume queues in which historical data has stationary patterns.	Low- to high-volume queues with seasonal patterns in a dimension other than time.
Data Required	A minimum of two years of history if seasonality is present; otherwise less data may be used. More data is helpful.	A minimum of two years of history if seasonality is present; otherwise less data may be used. More data is very advantageous.	A minimum of two years of history if seasonality is present; otherwise less data may be used. More data is helpful.	Several years of historical data to obtain good meaningful relationships.
When to Use	When the historical data has stationary patterns with or without trends or seasonality. When you can assume the past will continue to influence and represent the future.	When there is stable data that has regular correlations. When the patterns are very complicated with a combination of trend, seasonal, cyclical and random fluctuations.	When the historical data has stationary patterns with or without trends or seasonality. When you can assume the past will continue to influence and represent the future.	When cyclical patterns exist but the rise and fall is not within a fixed period. When the future would be based on something other than time.

** A "turning point" is where a trend or seasonal affect shifts direction*

	Time Series			Causal
	Weighted Moving Average	Box-Jenkins ARIMA	Exponential Smoothing	Multilinear Seasonal Regression
When to Use		Box-Jenkins performs better than Exponential Smoothing when the history is extended and more stable. Box-Jenkins is not well-suited for wildly variable, more volatile or short histories.	When a linear trend is present in the data.	
Constraints	May lag behind changes in trend if seasonal factors and other variables are not applied after the calculation. Complex relationships in data may be ignored.	These models can be very sophisticated and may be difficult to explain.	Can be less accurate for long term forecasts because it may lag behind changes in trends. Bad data (anomalies) in most recent periods can negatively impact the forecast.	Incomplete data or limited historical data can lead to false correlations. Increased complexity makes it difficult to interpret or determine the best model.
Captures Seasonality & Trends	No	Yes (based on pattern recognition)	Yes (based on structural analysis)	Yes (based on cause and effect relationship)

SUMMARY

An accurate forecast that optimizes schedules days, weeks or months in advance helps keep your customers satisfied, your employees engaged and your bottom line healthy.

For decades, more than two thousand customers and 2.7 million users of NICE WFM have benefited from its unique, industry-leading forecasting algorithm, which uses interval-specific weighted moving average methodologies with seasonality, week of month, rate of change and automated anomaly detection.

Now, with NICE WFM 7.0, we have added powerful new forecast models, including Box-Jenkins ARIMA, exponential smoothing and multilinear seasonal regression, supported by best pick with artificial intelligence. Together, these next-generation forecasting algorithms enable you to create more accurate forecasts that empower your team leaders with the insight they need for better decision-making.

NICE