ABL Digital Technologies

VOLUME FORECASTING WORKBENCH

A Practitioner's Guide

What the tool does · The six classical methods · Auto-selection by MAPE
Inputs explained · Outputs and the Excel workbook
A worked example end-to-end · Limits and tips

Companion to the Erlang, Shift Planning and Deal Sizing explainers.

VOLUME FORECASTING WORKBENCH

A Practitioner's Guide

What the tool does · The six classical methods · Auto-selection by MAPE
Inputs explained · Outputs and the Excel workbook
A worked example end-to-end · Limits and tips

Companion to the Erlang, Shift Planning and Deal Sizing explainers.

1 · What This Tool Does

The Volume Forecasting Workbench takes your historical contact volume — monthly totals as a minimum, daily data if you have it — and produces two forecasts: a long-term monthly view (up to 24 months ahead) and a short-term weekly view (up to 13 weeks) that includes the day-of-week breakdown.

It runs six classical time-series methods in parallel — Seasonal Naïve, Holt-Winters Additive, Holt-Winters Multiplicative, Auto-SARIMA, Classical Decomposition, and a Fourier-based decomposable model — scores each against a holdout window from your own history, and picks the best fit. You can also override the auto-pick if you have a strong reason to.

Every forecast comes with 80% and 95% confidence bands, computed by bootstrap from the residuals of the selected method (or analytically if you prefer the faster, Normal-distribution-based bands).

Everything runs in your browser. No data is sent anywhere.

> *This tool sits at the start of the WFM workflow: Forecast → Erlang staffing requirement → Shift assignment → Deal cost. The other three tabs on the Tools page handle the subsequent steps.*

2 · Before You Begin

Time series

A sequence of values measured at regular intervals over time. Monthly contact volume is a time series. Daily is a finer-grained one.

Trend

The long-run direction of the series — growth, decline, or flat. The forecast methods estimate this from the history and project it forward.

Seasonality

A repeating pattern. Contact-centre volume usually has two layers of seasonality: annual (December typically heavier than June for retail; July heavier than December for tax-related), and weekly (Monday heavier than Saturday for most operations). The methods in this tool capture both.

Holdout

The most recent N months of your history that are deliberately NOT used to fit the model. Instead, the model is fit on the older months, and then asked to forecast the holdout months. We then compare the forecast against the actual values to score accuracy. Default holdout is 6 months.

MAPE — Mean Absolute Percentage Error

The standard accuracy metric for forecasts. Computed as the average, across the holdout months, of |actual – forecast| ÷ actual. Expressed as a percentage. Lower is better. A MAPE of 5% means the forecast was on average within 5% of actual; 20%+ usually means the method is mismatched to the series.

Confidence bands

A range around the point forecast. The 80% band means "we expect the actual value to fall inside this range 80 times out of 100"; the 95% band is wider and more conservative. These are computed from the size of past forecast errors.

3 · The Six Methods

Each method makes different assumptions about how trend and seasonality combine. The right method depends on the shape of your history. Auto-selection picks for you, but it helps to know what each one is doing.

3.1 Seasonal Naïve

The simplest possible method. Forecast for next January = last January's value, adjusted for the average trend across the most recent year. Use as a baseline — if a fancier method cannot beat this, something is wrong with the data.

```
forecast[t+h] = value[t+h – 12] + average_trend
```

3.2 Holt-Winters Additive

Maintains three running estimates: level (current average), trend (how fast the level is moving), and 12 seasonal offsets (how much each month differs from the average). Updates all three with each new observation using smoothing parameters α, β, γ. Forecasts as level + h·trend + seasonal-for-that-month.

"Additive" because the seasonal pattern is added to the level — appropriate when the size of the seasonal swing is roughly constant across years.

3.3 Holt-Winters Multiplicative

Same idea as additive but the seasonal effect is a multiplier rather than an offset. Appropriate when the seasonal swing grows with the level — for example, if December is always 30% above the trend (not +5,000 contacts regardless of trend). This is the more common case for growing contact-centre volumes.

```
forecast[t+h] = (level + h·trend) × seasonal_factor[month]
```

3.4 Auto-SARIMA

The Box-Jenkins family — autoregressive integrated moving average, with a seasonal component. SARIMA(p,d,q)(P,D,Q)[s] is a model that captures both short-term auto-correlation and seasonal auto-correlation. This tool fits SARIMA(1,1,1)(1,1,1)[12] with grid search over the four AR/MA parameters — a standard configuration that handles most monthly contact-centre series.

More mathematically expressive than Holt-Winters but harder to interpret. When it wins on MAPE, it usually wins by capturing auto-correlation that Holt-Winters misses.

3.5 Classical Decomposition

Splits the series into three additive components — trend, seasonal, and residual — using a centred moving-average filter. Extrapolates the trend linearly and reapplies the seasonal pattern. Very transparent: you can show finance the trend line and the seasonal multipliers and they will understand the forecast.

3.6 Decomposable (Fourier)

Models the series as a linear combination of a trend term plus a small number of sine and cosine waves at annual frequencies. The underlying maths is the same that powers Facebook's Prophet framework. Robust to small gaps in history and slow shifts in seasonal shape. Useful when the seasonal pattern is itself drifting.

```
forecast[t] = a + b·t + Σ_k=1..K [α_k·sin(2πk·t/12) +
β_k·cos(2πk·t/12)]
```

4 · The Inputs — Every Field Explained

4.1 Historical Data

Two ways to load:

- Paste — two columns (date, volume) separated by tab, comma, or whitespace. The tool auto-detects the date format (YYYY-MM, MMM YYYY, MM/YYYY, full dates).
- Upload — .csv or .xlsx. First sheet, first two columns.

Granularity is auto-detected. If you provide daily dates (YYYY-MM-DD), the tool builds both a daily series (for the day-of-week pattern) and aggregates to monthly (for the long-term forecast). If you

provide only monthly dates, the long-term forecast still works; the short-term weekly view uses an even distribution across days within each month.

Minimum: 18 months. Recommended: 36 months. The auto-SARIMA and Fourier methods benefit from longer history.

4.2 Forecast Parameters

Long-term Horizon

How many months ahead to forecast. 12, 18, or 24. The 24-month view is the most useful for hiring planning and deal sizing; 12 is sufficient for an annual operating plan.

Short-term Horizon

How many weeks ahead to forecast at weekly granularity. 4, 8, or 13. The 13-week view is the typical roster-planning horizon.

Holdout (months)

Most recent N months held back from fitting and used for accuracy scoring. Default 6. Use more (e.g. 12) if you have lots of history and want a more robust accuracy estimate; less if your history is tight.

Annual Growth Override (%)

Optional. When set, multiplies the entire forecast by a compound annual factor. Useful for imposing a planning assumption — for example, "we know there's a contractual ramp of 8% per year that the history doesn't reflect yet." Leave blank to use the pure statistical forecast.

Method Override

Auto-select picks the lowest-MAPE method on holdout. Override if you have a reason — for example, if you want consistency with a previous forecast that used a specific method, or if the auto-pick looks visually wrong.

Confidence Band Method

Bootstrap (default) draws random samples from the residuals to build an empirical distribution of forecast error — non-parametric, realistic, slightly slower. Analytical assumes residuals are normally distributed and uses ±1.282 σ for 80% bands and ±1.96 σ for 95% — fast, but understates uncertainty when residuals are skewed.

4.3 India Holiday Calendar

A built-in list of India public holidays for 2024-2027 — Republic Day, Holi, Eid, Independence Day, Janmashtami, Gandhi Jayanti, Diwali, Christmas, and others. All enabled by default. You can:

- Toggle individual holidays off (right-click the checkbox to disable a specific date)
- Add custom holidays in the "YYYY-MM-DD Label" format

Active holidays inform the day-of-week pattern when daily history is provided. The decomposable Fourier model also picks up recurring holiday effects implicitly through its seasonal terms.

5 · The Outputs

5.1 Forecast Summary Card

- Selected Method — which of the six was picked (or your override)
- Holdout MAPE — the accuracy score of the selected method
- Total Forecast Volume across the long-term horizon
- Peak Forecast Month — the busiest month and its value
- Average Monthly Forecast
- Year-1 and Year-2 totals
- Year-on-Year Growth

5.2 Method Comparison Table

Every method, with its holdout MAPE, MAE (Mean Absolute Error), RMSE (Root Mean Square Error), and a ★ next to the selected one. Useful for understanding why one method won and whether the choice is sensitive — if multiple methods score within a percentage point of each other on MAPE, the auto-pick is close to arbitrary.

5.3 Long-term Chart

History (solid dark line) flows into forecast (dashed gold line) with the 80% and 95% confidence bands shown as shaded regions. Useful for visual sanity-checking — does the forecast continue the trend and seasonal shape of the history?

5.4 Short-term Chart

Weekly bars for the next 13 (or 4 or 8) weeks. When daily history was provided, the day-of-week pattern is baked into the weekly totals; the Excel download includes the daily breakdown.

5.5 Forecast Table Preview

First 12 months in a compact table: month, point forecast, 80% band, 95% band. Full horizon in the Excel.

6 · The Excel Workbook — Six Sheets

Clicking Download Excel produces ABL_Volume_Forecast.xlsx, styled consistently with the other workbooks on the Tools page.

6.1 Summary

Title, timestamp, INPUTS (every input echoed back), SELECTED METHOD with its accuracy metrics, and FORECAST HEADLINES (total volume, peak month, Year-1, Year-2, YoY growth).

6.2 Method Comparison

Every method with its MAPE, MAE, RMSE, Status (OK or Failed), and the ★ marker for the selected one. The Status column is useful when a method fails on short history — typically Auto-SARIMA needs at least 24 months to fit cleanly.

6.3 Long-Term Forecast

Every month in the horizon with Point Forecast, 80% Low, 80% High, 95% Low, 95% High.

6.4 Short-Term Forecast

Every week in the short-term horizon with the daily breakdown: Mon, Tue, Wed, Thu, Fri, Sat, Sun columns. Total column to the left.

6.5 Residuals

Last N months of history (the holdout) with Actual, Fitted, and Residual columns. Useful for spotting systematic over- or under-forecasting and for sharing with the analyst who will own the forecast.

6.6 Assumptions

A methodology narrative covering all seven points: the six methods, auto-selection, confidence band computation, short-term derivation, growth override, holiday calendar, and what the tool does not cover. Useful for the finance reviewer.

7 · A Worked Example

Suppose you have 36 months of monthly contact volume for a retail banking inbound queue, pasted into the input area:

```
2023-01 18,000
2023-02 16,500
2023-03 19,200
... ...
2025-12 24,800
```

Click Load & Validate Data. The status line confirms: "Loaded 36 months (2023-01 to 2025-12)".

7.1 Backtest

Set Long-term Horizon to 24 months, Holdout to 6, leave Method on Auto-select. Click Run Forecast. The tool fits all six methods on months 1–30 and forecasts months 31–36, comparing against actual:

```
Seasonal Naïve MAPE 11.2%
Holt-Winters Additive MAPE 6.8%
```

```
Holt-Winters Multiplicative MAPE 5.4% ★
Auto-SARIMA MAPE 5.9%
Classical Decomposition MAPE 7.1%
Decomposable (Fourier) MAPE 6.3%
```

Holt-Winters Multiplicative wins, presumably because the seasonal swing in the data is proportional to the level (December is always ~30% above the year, not always +6,000 contacts).

7.2 Forecast

The tool now refits Holt-Winters Multiplicative on all 36 months and forecasts 24 months ahead. The summary card shows:

- Selected Method: Holt-Winters Multiplicative
- Holdout MAPE: 5.4%
- Total Forecast Volume (24 months): ~580,000 contacts
- Peak Forecast Month: 2026-12 · 29,500
- Year-1 Total: ~278,000 · Year-2 Total: ~302,000 · YoY Growth: 8.6%

7.3 Confidence Bands

At 2026-12 the point forecast is 29,500. The 80% band might be 27,200–31,800; the 95% band 25,400–33,600. The wider the bands, the more uncertain the forecast — typically the bands widen as you go further into the future.

7.4 Short-term Detail

Because the example only provided monthly data, the short-term forecast distributes volume evenly across days within each month. If daily data had been provided, the 13-week view would show Monday-heavier and Sunday-lighter days reflecting the actual day-of-week pattern from history.

7.5 What You Do With This

Hand the long-term forecast to:

- The Shift Planning tool — to derive the weekly staffing requirement.
- The Deal Sizing tool — to derive the FTE roster and the cost-to-serve over the horizon.
- Finance — for the budget and hiring plan, with the residuals sheet as evidence of forecast quality.

8 · What This Tool Will Not Tell You

Things deliberately out of scope:

- AHT or handle-time forecasting (volume only). Forecast AHT separately if needed.
- Multi-channel allocation. Run the tool once per channel.

- Multi-skill or multi-site allocation of the forecast.
- ML methods (gradient boosting, LSTM, transformers). For typical contact-centre series these add complexity without much accuracy gain over the classical methods used here.
- Real-time or streaming re-forecasting. The tool is designed for periodic (typically monthly) re-runs as new history arrives.
- Outlier detection beyond what shows up in the residuals. If you know a month was anomalous (system outage, marketing spike), exclude it from history before pasting.

9 · Tips and Common Questions

How much history do I need?

Minimum 18 months. Below that the seasonal methods cannot estimate 12 seasonal offsets reliably. 24 months is workable but tight. 36+ months is comfortable and is what the auto-SARIMA and Fourier methods really want.

Why did my Auto-SARIMA fail?

Usually because the history is too short. With 24 months and d=1, D=1 differencing, you have only 11 effective observations left to fit the AR/MA parameters. Add more history or fall back to Holt-Winters Multiplicative, which works fine on 24 months.

The forecast looks too flat compared to my history

Check the trend in the most recent year against the long-run trend. If recent months have been growing faster than the older period, classical decomposition (linear-trend extrapolation) will under-forecast. Holt-Winters and SARIMA respond faster to recent changes — the auto-pick usually catches this. You can also use the Growth Override to impose your view.

Why do my confidence bands look enormous?

Either your history is noisy (high residual variance — try the "smooth" Holt-Winters methods instead of Naïve), or your forecast horizon is too long (bands widen with the square root of horizon). A 95% band at 24 months for a moderately-noisy series typically spans ±20%. If yours is much wider, the data or the method has an issue worth investigating.

Should I trust the auto-pick blindly?

Not blindly. Glance at the method comparison table — if the top three methods are within 1 MAPE point of each other, the choice is close to arbitrary and the human eye on the chart matters. Pick the one whose forecast looks most plausible to you.

Where can I learn more about the methods?

Hyndman & Athanasopoulos, "Forecasting: Principles and Practice" (free online at otexts.com/fpp3) is the standard reference and covers everything used here in depth. The Erlang Formulas Explainer (companion document) covers the queueing maths that consumes the forecast on the next step of the WFM workflow.