

narwhals

narwhals in brief, writing dataframe code once and running it on pandas, Polars, PyArrow, and more.

Wrap & Unwrap



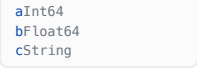



The bridge: turn any native frame into a narwhals frame and back.

PURPOSE	COMMAND	RESULT
Wrap a native frame.	<code>nw.from_native(df)</code>	
Restrict to eager frames.	<code>nw.from_native(df, eager_only=True)</code>	
Allow a Series through too.	<code>nw.from_native(s, allow_series=True)</code>	
Hand the result back.	<code>nw.to_native(df)</code>	
Build a frame from scratch.	<code>nw.from_dict(data, backend="polars")</code>	

narwhals borrows your frame; to_native returns the exact same native type

Inspect & Metadata

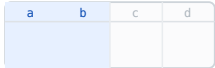
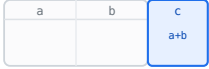

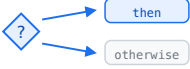


Look at a frame the same way no matter what backend it wraps.

PURPOSE	COMMAND	RESULT
Get (rows, cols).	<code>df.shape</code>	
List the column names.	<code>df.columns</code>	
Name to dtype mapping.	<code>df.schema</code>	
Which engine is underneath?	<code>df.implementation</code>	
Reach the native module.	<code>nw.get_native_namespace(df)</code>	
Convert to a concrete backend.	<code>df.to_pandas() / .to_polars()</code>	

every frame answers shape/columns/schema the same, whatever is underneath

Select, Columns & Expressions

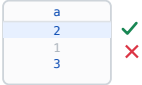
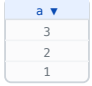
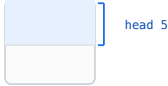

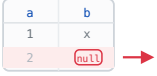
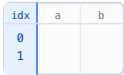
Build columns with `nw.col(...)` expressions, Polars-style.

PURPOSE	COMMAND	RESULT
Pick columns.	<code>df.select(nw.col("a", "b"))</code>	
Add / overwrite columns.	<code>df.with_columns((nw.col("a") + nw.col("b")).alias("c"))</code>	
Reference a literal value.	<code>nw.lit(0)</code>	
Conditional column.	<code>nw.when(nw.col("a") > 1).then(..).otherwise(..)</code>	
Combine columns horizontally.	<code>nw.sum_horizontal("a", "b")</code>	
Pick columns by selector.	<code>df.select(ncs.numeric())</code>	

`nw.col(...)` expressions are reusable recipes, run by `select` / `with_columns`

Filter, Sort & Transform Rows

Row-level operations: keep, order, dedupe, reshape.

PURPOSE	COMMAND	RESULT
Keep rows matching a test.	<code>df.filter(nw.col("a") > 1)</code>	
Sort by one or more columns.	<code>df.sort("a", descending=True)</code>	
First / last N rows.	<code>df.head(5) / .tail(5)</code>	
Drop duplicate rows.	<code>df.unique(subset=["a"])</code>	
Drop missing values.	<code>df.drop_nulls()</code>	
Add a row index.	<code>df.with_row_index()</code>	

these verbs read the same on a huge Polars LazyFrame and a tiny pandas frame

Group, Aggregate & Join

Split-apply-combine and relational joins, backend-agnostic.

PURPOSE	COMMAND	RESULT
Group then aggregate.	<code>df.groupby("k").agg(nw.col("v").sum())</code>	
Multiple aggregations.	<code>.agg(nw.col("v").mean(), nw.len())</code>	
Inner / left / outer join.	<code>df.join(other, on="id", how="left")</code>	
Cross join.	<code>df.join(other, how="cross")</code>	
Stack frames vertically.	<code>nw.concat([df1, df2], how="vertical")</code>	
Window expression over groups.	<code>nw.col("v").sum().over("k")</code>	

narwhals maps each verb to the backend's own optimized implementation

Expression Namespaces

Typed sub-APIs (`str` / `dt` / `cat`) that work the same on every backend.

PURPOSE	COMMAND	RESULT
Uppercase strings.	<code>nw.col("name").str.to_uppercase()</code>	
Test & slice text.	<code>nw.col("name").str.contains("a")</code>	
Extract date parts.	<code>nw.col("ts").dt.year()</code>	
Truncate datetimes.	<code>nw.col("ts").dt.truncate("1d")</code>	
Categorical to categories.	<code>nw.col("c").cat.get_categories()</code>	
Parse strings to dates.	<code>nw.col("d").str.to_datetime(format="%Y-%m-%d")</code>	

`.str` `.dt` `.cat` `.list` `.struct` make typed work uniform across engines

Lazy Frames & I/O

Defer work with LazyFrame; read and write without committing to one engine.

PURPOSE	COMMAND	RESULT
Go lazy (defer execution).	<code>df.lazy()</code>	
Trigger the computation.	<code>lf.collect()</code>	
Scan a CSV lazily.	<code>nw.scan_csv("data.csv", backend="polars")</code>	
Read a CSV eagerly.	<code>nw.read_csv("data.csv", backend="pandas")</code>	
Stream a lazy result to disk.	<code>lf.sink_parquet("out.parquet")</code>	
Write a frame out.	<code>df.write_parquet("out.parquet")</code>	

`read_*` and `scan_*` require `backend=`; narwhals will not guess the engine

Write Portable Functions

The payoff: one function, every dataframe library.

PURPOSE	COMMAND	RESULT
Auto-wrap/unwrap a function.	<code>@nw.narwhalify def f(df): ...</code>	
Write the body once.	<code>return df.with_columns(nw.col("a").mean())</code>	
Manual wrap form.	<code>df = nw.from_native(df); ... return nw.to_native(df)</code>	
Use the stable API.	<code>import narwhals.stable.v1 as nw</code>	
Pass native through untouched.	<code>nw.from_native(x, pass_through=True)</code>	
Convert a Series the same way.	<code>nw.from_native(s, series_only=True)</code>	

`@nw.narwhalify` accepts and returns whatever dataframe library the caller uses