



On-Demand Training: Additional Filtering Topics Transcript

Welcome to this video on additional filtering topics. You can download the Exercises workbook to follow along in your own copy of Tableau.

Filtering in Tableau is incredibly powerful, but with that it can get complicated. This video assumes a basic understanding of filtering as laid out in the previous videos and goes beyond into some deeper concepts.

Context Filters

First is Context Filters.

- One thing to note about the filter shelf is that the filters are computed independently of each other.
- If a given filter will trim the dataset down, it may be worth Adding to Context.
- A Context Filter will be computed first, then all other filters will run only on its results.

For example, if we only wanted to look at products within the Machines sub-category, we may want to make that a context filter by right clicking and selecting “Add to Context”.

- The filter is now grey, and any subsequent filters will only be on machine products.
- Context filters are generally best to use when they’ll return a significantly smaller set of results, the guideline is about a tenth or less of the original data
- Context filters ideally shouldn’t be something that would change frequently, so let’s remove this quick filter

Apply to Worksheets

When a pill is placed on the filter shelf, by default that filter applies only to the current worksheet. If it makes sense to have the filter applied more broadly, we can change the scope by clicking on the pill’s dropdown to bring up the menu, “Apply to Worksheets” > “All Using this Data Source” or “Selected Worksheets”.

This can be especially useful on a Dashboard, where a filter that is relevant to multiple views can be set to apply to all or some of them simultaneously.

- By default, the quick filter only applies to the view it was brought out for
- But if we go to the menu and say “apply to worksheets” “all using this data source”
- Now all the views filter

And not just the views on the dashboard.

- Let’s leave it on this category and build a new sheet.
- Double click Sales, double click Category
- Note, we only have one Category
- But fortunately, we get this visual indication here that a filter has been applied.
 - This icon lets you know there’s a filter on this data connection, so we can backtrack and hunt it down if that’s not what we want
- Let’s leave this Quick Filter on All

- But another solution would be to say “Apply to worksheets” “selected worksheets” and target just the ones we want to have filter.
- It’s important to note that taking an “applied to all” filter off the filter shelf removes it from all views, not just the one we’re looking at, so be careful when using this option

Which Field to Filter On

Because of the power to control exactly what is filtered on and how, it’s important to recognize which field needs to be filtered and in what way. Let’s say we wanted to run a promotion for customers who purchased more than \$10,000 in sales.

- A straightforward way to do that might be to build a view of customers by sales, and filter for Sales, sum is at least \$10,000.
- However, this summary filter of sales >\$10,000 is dependent on the structure of the view
- If we modify the view, we risk altering how the filter is applied
 - Here, we have 621 rows, indicating 621 customers eligible for our promotion
 - But if we break it out by Color on Category, suddenly we’re down to 81 rows. Why?
 - That many customers didn’t have over 10,000 in sales in a single category, even if they may have been eligible overall.
 - Tableau ran each mark past the filter – and there’s a mark per customer per category in this view.

Alternatively, we could bring Customers to the filter shelf, and use the condition tab to say we only want customers who have sales greater than 10 000

- We have the same initial 621 rows, and if we again break it out by Category
- We still have 621 rows. No customers were filtered out like they were last time.
 - Here, Tableau ran Customers by their sum of sales through the filter, then any who had more than \$10,000 were put into the view, regardless of how the Categories split up those sales

Knowing how we set up our filter will impact how we can manipulate the view and maintain the filter as we designed it.

Filter Order of Operations

What order filters are put onto the filter shelf usually has no impact on the query sent to the data source, filters are grouped together into the WHERE clause of the SQL query. Exceptions include:

- Context Filters, which are applied first and all subsequent filters are run against their temporary table of results
- And Table Calculation filters, which are applied last (because these are performed only on the data in the view)

Performance and Filter Queries

Inefficient filtering is one of the most common causes of poorly performing workbooks and dashboards. For detailed information, see the related materials below this video.

Some things to consider about filtering and its impact on performance:

- Context filters are slow to create or change due to the temporary table being created behind the scenes, but once they are in place they can increase performance since further queries are only run off a subset the subset of data as opposed to the entire data set
- Filtering large lists of discrete options can be slow
- Cascading quick filters can help narrow down the list of options, however:
- When a filter requires Tableau to find all potential field values, this requires a complex query that can take time to evaluate.
 - (things such as use “only relevant values”, or quick filters layouts of multiple or single value lists)
 - Or when “Keep Only” and “Exclude” involve a large number of marks, Exclude can be faster than Include

The presence and maintenance of indexes in the data source can dramatically improve filtering performance

Conclusion

Thank you for watching these additional filtering topics. We invite you to continue with the On Demand Training videos to learn more about using Tableau Desktop