



Alan Eldridge and Tara Walker

---

# Understanding Level of Detail Expressions

At Tableau, our goal is to make data analysis a delightful experience. People tell us that when they are deeply engaged in Tableau they stop thinking about the mechanics of using the product and just have fun asking questions of their data. We call this experience flow—a state of joyful immersion in a task.

If you have to start thinking about how to use the tool to solve the problem, the state of flow is broken. One common cause of this is the need to work with data that has been aggregated to different levels of detail. These questions are often simple to ask, but hard to answer. They often sound something like:


- Can I plot the number of days per quarter where my company had more than 100 orders?
- Find the biggest deal each sales person has ever closed, then show the averages by manager?
- Tag every customer by the year he/she first became a customer, then use that tag to group sales?

To address these types of questions, Tableau 9.0 introduces a new syntax called Level of Detail (LOD) Expressions. This new syntax both simplifies and extends Tableau's calculation language by making it possible to address level of detail questions directly.

**LOD Expressions represent an elegant and powerful way to answer questions involving multiple levels of granularity in a single visualization.**

# Table of Contents


How LOD Expressions Work—Explaining The “Level Of Detail” .....	4
Include: Calculating At A Lower Level Of Detail .....	7
Exclude: Calculating At A Higher Level Of Detail.....	12
Fixed: Specifying The Exact Level Of Detail.....	14
Canonical Use Cases.....	18
Final Thoughts.....	21
Appendix A – Additional Resources .....	21
Appendix B – LOD Expressions vs. Table Calculations .....	22




## How LOD Expressions Work— Explaining The “Level Of Detail”

A key aspect of exploring data is understanding the structure of the source. For example, you may have restaurant inspection data that at the most granular level is listed by its street address. You may then want to aggregate the data to view properties by zip code, city, state, or even country.

In Tableau, you typically do this by dropping the dimensions you care about into your view (e.g., city, state). Depending on the dimensions you've chosen to add to the view, your data will be aggregated accordingly —to the “viz level of detail”, or Viz LOD for short.




*Placing dimensions on the shelves highlighted here will add them to the viz LOD.*



*Placing dimensions on the pages, filters and tooltips shelves does not add them to the Viz LOD. Dimensions placed on these shelves are ways of modifying the data in the view without displaying it visually.*

Now here is a key point: LOD Expressions allow you to determine the levels of detail (i.e., the dimensions) used in a calculation without actually dropping those dimensions into the visualization. You can define the level of detail at which a calculation should be performed, independently of the Viz LOD.

In the following example (using restaurant inspection data), two dimensions have been added to the view: City and State.




*This map shows restaurant inspection data from YELP in the greater Edinburgh area. The data has been aggregated to the City/State level of detail.*

The data in the view is aggregated based on the Viz LOD—which in this case consist of City and State—and is more aggregated than the underlying data source. The selected point in the image shows the average user fans for all restaurants in Newbridge, Edinburgh.


Adding more granular dimensions to the view will result in a less aggregated Viz LOD. For instance, we could add Business ID to the visualization (by dragging it on the Detail Shelf), to see the average user fans for every individual business. Doing so will also change the visualization: every single business will appear as a circle on the map. But what if we don't want the visualization to change? What if we want to determine the total user fans for each Business ID, average those values for each city, and show only one circle per city? What we want to see is the average number of fans per restaurant in each city.

This will require adding an additional dimension to the view without dragging that dimension into the visualization. An LOD expression will allow us to this.

Let's create a new calculated field called Fans per Business. Here is a brief introduction to the syntax:



The expression tells Tableau to perform the aggregation for each Business ID, regardless of other dimensions used in the viz. You can use the LOD Expression to calculate the total User Fans per Business ID. After dragging this new field into the view, we can then average those values per city.



*The Fans per Business field has been added to the Color Shelf. Newbridge has the highest average fans per business with 185 fans—a value that was computed using a FIXED LOD expression.*

By using the FIXED operator in our LOD Expression, we gain insight into which cities have, on average, more fans per Business ID. Meaning, those cities with a darker shade of blue have more popular restaurants (or the city could have more residents and hence, more total fans per restaurant).

**There are three types of LOD expression keywords—EXCLUDE, INCLUDE and FIXED—each of which alters the scope of the LOD expression.**

## Include: Calculating At A Lower Level Of Detail

In this example, we'll look at a standard sales database (the Superstore database that comes with Tableau). Here, each row represents the sale of a certain item. Order may contain multiple items and spread across multiple rows. In other words, the deepest level of detail in this database is a unique item.

View Data: Sample – Superstore							
Customer Name	Order Date	Order ID	Sub-Category	Category	Product Name	Quantity	Sales
Claire Gute	11/9/13	CA-2013-152156	Bookcases	Furniture	Bush Somerset Collection Bookcase	2	\$262
Claire Gute	11/9/13	CA-2013-152156	Chairs	Furniture	Hon Deluxe Fabric Upholstered Sta...	3	\$732
Darrin Van Huff	6/13/13	CA-2013-138688	Labels	Office Supplies	Self-Adhesive Address Labels for T...	2	\$15
Sean O'Donnell	10/11/12	US-2012-108966	Tables	Furniture	Bretford CR4500 Series Slim Recta...	5	\$958
Sean O'Donnell	10/11/12	US-2012-108966	Storage	Office Supplies	Eldon Fold 'N Roll Cart System	2	\$22
Brosina Hoffman	6/9/11	CA-2011-115812	Furnishings	Furniture	Eldon Expressions Wood and Plasti...	7	\$49
Brosina Hoffman	6/9/11	CA-2011-115812	Art	Office Supplies	Newell 322	4	\$7
Brosina Hoffman	6/9/11	CA-2011-115812	Phones	Technology	Mitel 5320 IP Phone VoIP phone	6	\$907
Brosina Hoffman	6/9/11	CA-2011-115812	Binders	Office Supplies	DXL Angle-View Binders with Locki...	3	\$19
Brosina Hoffman	6/9/11	CA-2011-115812	Appliances	Office Supplies	Belkin FSC206VTEL 6 Outlet Surge	5	\$115
Brosina Hoffman	6/9/11	CA-2011-115812	Tables	Furniture	Chromcraft Rectangular Conferenc...	9	\$1,706
Brosina Hoffman	6/9/11	CA-2011-115812	Phones	Technology	Konfet 250 Conference phone - C...	4	\$911
Andrew Allen	4/16/14	CA-2014-114412	Paper	Office Supplies	Xerox 1967	3	\$16
Irene Maddox	12/6/13	CA-2013-161389	Binders	Office Supplies	Followes PB200 Plastic Comb Bindi...	3	\$408
Harold Pawlan	11/22/12	US-2012-118983	Appliances	Office Supplies	Holmes Replacement Filter for HEP...	5	\$69
Harold Pawlan	11/22/12	US-2012-118983	Binders	Office Supplies	Storex DuraTech Recycled Plastic F...	3	\$3
Pete Kriz	11/11/11	CA-2011-105893	Storage	Office Supplies	Stur-D-Stor Shelving, Vertical 5-S...	6	\$666
Alejandro Grove	5/13/11	CA-2011-167164	Storage	Office Supplies	Followes Super Stor/Drawer	2	\$56
Zuschuss Donatelli	8/27/11	CA-2011-143336	Art	Office Supplies	Newell 341	2	\$9
Zuschuss Donatelli	8/27/11	CA-2011-143336	Phones	Technology	Cisco SPA 501G IP Phone	3	\$213
Zuschuss Donatelli	8/27/11	CA-2011-143336	Binders	Office Supplies	Wilson Jones Hanging View Binder...	4	\$23
Ken Black	12/10/13	CA-2013-137330	Art	Office Supplies	Newell 318	7	\$19
Ken Black	12/10/13	CA-2013-137330	Appliances	Office Supplies	Acco Six-Outlet Power Strip, 4' Cor...	7	\$60
Sandra Flanagan	7/17/14	US-2014-156909	Chairs	Furniture	Global Deluxe Stacking Chair; Gray	2	\$71
Emily Burns	9/25/12	CA-2012-106320	Tables	Furniture	Bretford CR4500 Series Slim Recta...	3	\$1,045
Eric Hoffmann	1/16/13	CA-2013-121755	Binders	Office Supplies	Wilson Jones Active Use Binders	2	\$12
Eric Hoffmann	1/16/13	CA-2013-121755	Accessories	Technology	Imation RCR Mini TravelDrive 11CR	2	\$01


A snapshot of the Superstore database that comes with Tableau.

The first row in this snapshot of the database is for a purchase of 2 Bush Somerset Bookcases. The second row is a purchase of 3 Hon Stacking Chairs. These two rows together comprise a single order—namely, order CA-2013-152156.


Suppose you are analyzing the sales performance of each region and would like to know which region has the highest (or lowest) average order size?

To figure this out, you need to calculate the size of each order (sum the sales corresponding to each Order ID), and then average those values by region.

This business question is easy enough to ask, and with the new LOD Expressions syntax, Tableau makes it easy to answer. Here is a more detailed discussion of the new syntax:



LOD Expressions can be written in the calculation editor as shown here. This LOD expression is used to sum the purchases in each Order ID. The result is a new field called **Order Size**.



The bars on the left show the average size of orders by Region—computed using an LOD Expression—while the bars on the right show the average Sales by Region (i.e., the average of all order line items, regardless of which order they were in). You may now answer the question: Which Region has the highest average order size?

You can see that North Asia and Central Asia have the highest average Order Size, \$737 and \$733 respectively. We were able to determine this even though Order ID does not appear in the visualization. (Before Tableau Version 9, we could not have computed these values unless Order ID had been added to the view).

If we had simply plotted Region versus AVG(Sales)—as seen in the bars on the right in the figure—we would see the average of all line items (rows) in each Region, which is not the result we are looking for. In contrast, with the LOD expression Order Size, we are able to first determine the size of each order (i.e. the sum of sales for all line items within that order), and then average the resulting orders by Region to determine average Order Size by Region.

Now that we have a sense of our largest average order size, let's ask a slightly more complex question: Which country in the sales database has sales reps who close the “biggest deals,” on average? What we want to do is:

1. Find the biggest deal (the max deal) that each sales person has ever closed, and then
2. Average those ‘biggest deals’ by country.

This question is multi-faceted, but it is easy to answer with an LOD expression:

The screenshot shows the Tableau Calculation Editor window. The title bar says "Avg Largest Sales Deal by Rep". The main text area contains the LOD expression: `AVG({INCLUDE [Sales Rep]: MAX([Sales])})`. Below the editor, a message says "The calculation is valid." There are "Apply" and "OK" buttons at the bottom right. A small note above the "OK" button says "Sheets Affected ▾".

*The LOD expression Avg Largest Sales Deal by Rep is used to calculate the average maximum deal per sales rep. In this case, the average of the LOD Expression is typed directly in the calculation editor window.*



The LOD Expression Avg Largest Sales Deal by Rep allows us to compare the average max deal across Countries. The average “biggest deal” for all sales reps in Tajikistan is \$4,533.


Notice that we answered this question with one expression AVG({INCLUDE [Sales Rep] : MAX([Sales])})—no need for complex formulas. In fact, you can ask additional questions of your data by adding additional dimensions to the view, and the calculation will update. For instance, let’s add Year to the analysis: “Our clients are amazed and very impressed with the richness of the way we can present information visually,” he says. “It helps them in their decision-making.”



The LOD expression Avg Largest Sales Deal by Rep will update when additional dimensions are added to the view—in this case Year of Transaction Date. The visualization now shows the average “biggest deal” each sales rep closed, by country and year. The average of the “biggest deals” for all sales reps in Tajikistan was \$1,636 in 2012, \$3,482 in 2013, and \$2,251 in 2014.

Using the INCLUDE keyword in the calculation, the Sales Rep field is being explicitly included in the calculation, but so are any other dimensions that are placed in the visualization (in this case Country and Year). By adding Year to the view, we can dive deeper into our analysis and can now gain insights such as this: Bahrain had the largest average “biggest deal” in 2012 with \$4,069.

A graphical representation of how Tableau performs an INCLUDE LOD Expression is depicted in the following flow diagram.



The INCLUDE keyword creates an expression that is less aggregated (i.e., more granular) than the Viz LOD. The specified dimension(s) are first added to the viz LOD before calculations are performed.


Notice the INCLUDE Expression is used in the view as an aggregated measure. In fact, all INCLUDE expressions are either used as measures or aggregated measures when placed on the view.

## Exclude: Calculating At A Higher Level Of Detail

Consider the following scenario: For each month, we want to see the total Sales, as well as the total sales by Region. To do this:

1. We need to exclude Region from our calculation of the monthly Total Sales
2. And then include Region when calculating the regional Sales breakdowns.

Let's explore an additional example using the sales database—as described earlier.



Monthly total Sales (color on bars) is computed using an EXCLUDE LOD expression. The sum of Sales by Region (the length of bars) is not based on an LOD expression. The result is a single visualization showing numbers at two different levels of detail.

**Total Sales**

---

{EXCLUDE [Region] : SUM( [Sales] )}

---


The calculation is valid.

This LOD expression, called Total Sales, allows you to compute the monthly total sales across all Regions.

Notice that in the viz above, Region has been placed on the Column shelf and thus contributes to the Viz LOD of Region, Month(Order Date). Using the EXCLUDE expression, you have the ability to calculate the total sales (across all regions) while displaying the regional sales breakdowns. As such, we have created an LOD expression that is “above” the Viz LOD (i.e. less granular):

**{FIXED [Business ID] : SUM([User Fans])}**

A key to the EXCLUDE keyword: Tableau first removes the excluded dimension from the Viz LOD and performs the calculation as if the dimension was not present at all. The result is then displayed visually. A graphical representation of how Tableau performs an EXCLUDE LOD Expression is depicted in the following flow diagram.




Using an EXCLUDE LOD expression will exclude the desired dimension(s) from the calculation.

The expression {EXCLUDE [Region] : SUM([Sales])} tells Tableau to calculate the sum of sales using whatever dimensions are in use in the viz, excluding the Region dimension. This means we end up with a single value for each month which represents the total sales across all regions.

Now we have a powerful view that shows us both total Sales and Sales by Region—both of which use the SUM aggregation.

We can also mix aggregations. For instance, let's change our LOD Expression to show the average Sales per month, while still showing regional sums:



*The color now represents the average monthly Sales and is computed using an EXCLUDE LOD expression. The sum of Sales by Region (the length of the bar) is not based on an LOD expression.*

Similar to an INCLUDE Expression, all EXCLUDE Expressions are either used as measures or aggregated measures when placed on the view. These types of expressions are great for “percent of total” or “difference from overall average” calculations.

## Fixed: Specifying The Exact Level Of Detail

LOD Expressions also open the door to creating an aggregation level completely independent of the Viz LOD; something that was previously only possible by using custom SQL.

As an example, consider the following scenario: Let's say you'd like to analyze YELP data to find the yearly cohorts in which a business had its first review. Does each cohort have the same review trends?

With LOD Expressions, we can specify the cohort at an exact level of detail:




*This LOD expression will fix the level of detail to each Business ID. Then, for each Business ID, it takes all reviews and finds the minimum Year of the Review Date and associates that value with the Business ID. You can imagine First Review Year to be a new column in the database.*

When we use this field in the viz, the calculation scope is implicitly defined in the expression. Once the First Review Year is recorded for each Business ID (as shown below), exploring the cohorts can be insightful.

<b>Business - Id</b>	<b>Year of Review - Date</b>	<b>First Review Year</b>
<b>0_0XEqWPy4LizKqzC9fc8A</b>	2010	2010
	2011	2010
	2012	2010
	2013	2010
<b>0_1nEveztDCul8R2CxLK0A</b>	2012	2012
	2014	2012
<b>0_3hrGrrP-J5eeHRwQfZ6w</b>	2008	2008
	2009	2008
	2010	2008
	2011	2008
	2012	2008
	2013	2008
	2014	2008
<b>0_5QEcn3WbFvLQbhs_xWIQ</b>	2013	2013
	2014	2013


*The level of detail has been fixed at Business – ID. Notice the first review year (as found in the FIXED LOD Expression) is the same for each review row.*



This visualization shows the Number of Records (number of reviews) broken down by First Review Year and the Year of Review Date. By adding First Review Year to Color, we have created a cohort analysis.


The FIXED keyword allows us to specifically define a calculation's aggregation level. Unlike INCLUDE and EXCLUDE, it does this independently of the dimensions used in the visualization.

The result of a FIXED expression can be coarser or finer than the Viz LOD depending on the relationship between the FIXED dimension and the Viz LOD.



Depending on the granularity level of the Viz LOD in comparison to the fixed dimension, the resulting value could act similarly to an EXCLUDE expression (diagram on the left) or to an INCLUDE expression (diagram on the right).


Now, let's use First Review Year as a filter:



An LOD expression can also be used as a filter. By using a Quick Filter on First Review Year, you can quickly step through the yearly cohorts to see if certain areas of Las Vegas were first reviewed earlier or later than others. The FIXED LOD Expression is calculated above the Dimension filter on Business – State in the view.

Notice that each yearly cohort has been given a discrete “bin”—that is, First Review Year is used in the view as a dimension. FIXED expressions can be used as dimensions or measures. Depending on the data type, Tableau will place the resulting calculation either as a dimension or a measure.

A key distinction between INCLUDE/EXCLUDE and FIXED is where each falls in the filtering hierarchy as shown below. FIXED LOD Expressions are computed before dimension filters and after context filters. This can enable many use cases




*FIXED LOD expression filters are applied after context filters and before dimension filters. On the other hand, INCLUDE/EXCLUDE filters are applied after dimension filters and before measure filters.*


## Canonical Use Cases

A few useful scenarios for LOD Expressions have been outlined in previous sections of this paper—however, these are only the beginning of the power of using an LOD Expression to solve a business question. Other notable examples:


**Histogramm of Number of Orders:** How many customers in each segment have made 1, 2, 3 etc. orders?




**Actual vs Target:** What percentage of products are meeting their profit target per state?



**Customer Acquisition:** What is the cumulative customer acquisition by market?



**Daily Profit KPI:** How many days each month are highly profitable, profitable, or unprofitable?



For a thorough discussion of some of the best use cases for LOD Expressions, please refer to the Tableau publication “Top 15 LOD Expressions”, which includes online sample workbooks with step-by-step instructions.

## Final Thoughts

LOD Expressions are a powerful new capability of Tableau 9.0 that allow us to easily solve problems that previously required complicated formulas. They allow us to intuitively define the scope of calculations and stay in a state of flow as we explore our data.

LOD Expressions are not a new form of Table Calculations; they can replace many Table Calculations, but also open new possibilities. LOD Expressions and Table Calculations operate differently. A Table Calculation is generated exclusively from the result of a query, while an LOD Expression is usually generated as part of the query to the underlying data. Table Calculations always produce measures as their result, while LOD Expressions can create measures, aggregated measures or dimensions.

LOD expressions represent a vital step towards the goal of complete flow, where all questions are simple and elegant to answer.

## Appendix A – Additional Resources

For additional resources on creating, using, and explaining LOD Expressions, please refer to the materials listed below.

- First Look: An introduction to [LOD Expressions](#) from Robin Cottiss
- Explanation: [Level of Detail Expressions](#) from Bora Beran
- Example and Tutorial: [Top 15 LOD Expressions](#)
- Example: Craig Bloodworth of the Information Lab [records his experience](#) using LOD expressions

## Appendix B – LOD Expressions vs. Table Calculations

For people who used Tableau before Version 9, LOD Expressions will allow you to replace calculations you may have authored in more cumbersome ways:

- Using Table Calcs, you may have tried to find the first or last period within a partition. For example, calculating the headcount of an organization on the first day of every month.
- Using Table Calcs, Calculated Fields, and Reference Lines, you may have tried to bin aggregate fields. For example, finding the average of a distinct count of customers.
- Using Data Blending, you may have tried to achieve relative date filtering relative to the maximum date in the data. For example, if your data is refreshed on a weekly basis, computing the year to date totals according to the maximum date.


This appendix is designed for people who first learned Tableau using Version 8 or an earlier release. In those versions of the software, Table Calculations were sometimes used to specify the level of aggregation. Although LOD Expressions now make it easy to pinpoint the level of detail, Table Calcs have not lost their important place in your analysis.

The table below outlines a few key distinctions between Table Calcs and LOD Expressions.

Table Calculations	LOD Expressions
Table calculations are generated from the query results.  WINDOW_AVG(SUM([Sales]))	LOD Expressions are generated as part of the query to the underlying data source. They are expressed as a nested select so are dependent on DBMS performance:  SELECT T1.[State],SUM(T2.[Sales per County, State])  FROM [DB Table] T1 INNER JOIN (SELECT [State], [County], AVG([Sales]) AS [Sales per County, State]  FROM [DB Table] GROUP BY [State],[County]) T2  ON T1.[State] = T2.[State]  GROUP BY T1.[State]
Table calculations can only produce results that are equal to or less granular than the viz LOD.	LOD Expressions can produce results that are more granular than, less granular than or totally independent of the viz LOD.
Dimensions that control the operation of a table calculation are separate from the calculation syntax and specified in the “compute using” or “running along” menus.	Dimensions that control the calculation for a LOD expression are embedded in the calculation syntax.  For INCLUDE/EXCLUDE expressions these are relative to the viz LOD, for FIXED expressions these are absolute.
	LOD Expressions can be used as measures, aggregated measures or dimensions. The data type is determined by the formula.  INCLUDE/EXCLUDE expressions are usually used as measures or aggregated measures. FIXED expressions are usually used as dimensions, but can also be used as measures.
Table calculations are always aggregated measures.	LOD Expressions can be used in other constructs.
Filters on table calculations act as a hide – they do not remove records from the result set.	Filters on LOD Expressions act as an exclude – they remove records from the result set. Depending on whether the expression is FIXED or INCLUDE/EXCLUDE it is evaluated at different stages of the visualisation pipeline.

When attempting to achieve the equivalent of an EXCLUDE statement—as seen earlier—a table calc may be difficult to use in a non-convoluted way. You’d need to address all dimensions in the viz as either partitioning or addressing—according to the desired level of detail.


Additionally, because table calculations are aggregated from the query results they can only work “up” from the viz LOD (i.e. more aggregated/less granular).



*Using table calculations as a method of specifying the level of detail relies on the table calc level of detail (i.e. setting dimensions as partitioning or addressing).*

The addressing fields are the dimensions that you'd like to exclude in the calculation.

On the other hand, if you'd like to achieve the equivalent of an INCLUDE statement using table calcs, you need to make the query results less aggregated to match the lowest level calculation.



To achieve the same result as an INCLUDE LOD expression prior to Tableau 9.0, you could utilize a table calc to create a new viz LOD.


The next step would be to use table calculations to aggregate back up to the original viz LOD. However there will be more viz LOD records than we need so it's generally necessary to keep only one using a filter like INDEX()=1 or an equivalent.

As you can see, using a table calc to achieve an EXCLUDE/INCLUDE expression technically arrives at the answer, but could be performed much easier (and quicker) with an LOD expression.

## About Tableau


Tableau helps people see and understand data. Tableau helps anyone quickly analyze, visualize and share information. More than 26,000 customer accounts get rapid results with Tableau in the office and on-the-go. And tens of thousands of people use Tableau Public to share data in their blogs and websites. See how Tableau can help you by downloading the free trial at [www.tableau.com/trial](http://www.tableau.com/trial).

### Additional Resources



[Download Free Trial](#)

### Related Whitepapers




[Why Business Analytics in the Cloud?](#)

[5 Best Practices for Creating Effective Campaign Dashboards](#)

[See All Whitepapers](#)

### Explore Other Resources



· [Product Demo](#)

· [Training & Tutorials](#)

· [Community & Support](#)

· [Customer Stories](#)

· [Solutions](#)

