

# Release Notes

We are happy to inform you that the latest version 4.4 of ThoughtSpot is available. Below is the list features and issues that were part of this release.

## New Features and Functionality

### Description

#### [Learn ThoughtSpot in BETA](#)

The Learn option from the main menu is a BETA feature that uses replay to teach you more about ThoughtSpot.

#### [Support for UK and Canadian Languages](#)

Support for UK English (**en-GB**) and Canadian English (**en-CA**) was added. Users can [set these locale's via their profile](#)

#### [Show regression line option](#)

Users have a new option to **Show Regression Line** on line, column, bubble, and scatter charts.



#### [New Has administration privileges](#)

If you add the privilege **Has administration privileges** to a group, note that all users in that group will be able to see all the data in ThoughtSpot. Administrators can see all data sources, and row-level security (RLS) does not apply to them. If you want to enable users to bypass row-level security, add them to a group with **Can Administer RLS** privi-

## Description

leges.

### *Check cluster command added*

The `tscli cluster check` command was added to check the status of all or specific nodes in a cluster.

### *Service logs stored to disk*

Service logs are now stored to disk. This ensures better protection against service logs from exceeding the SSD on a node and impacting a node.

### *The tscli nas mount-cifs command requires a password*

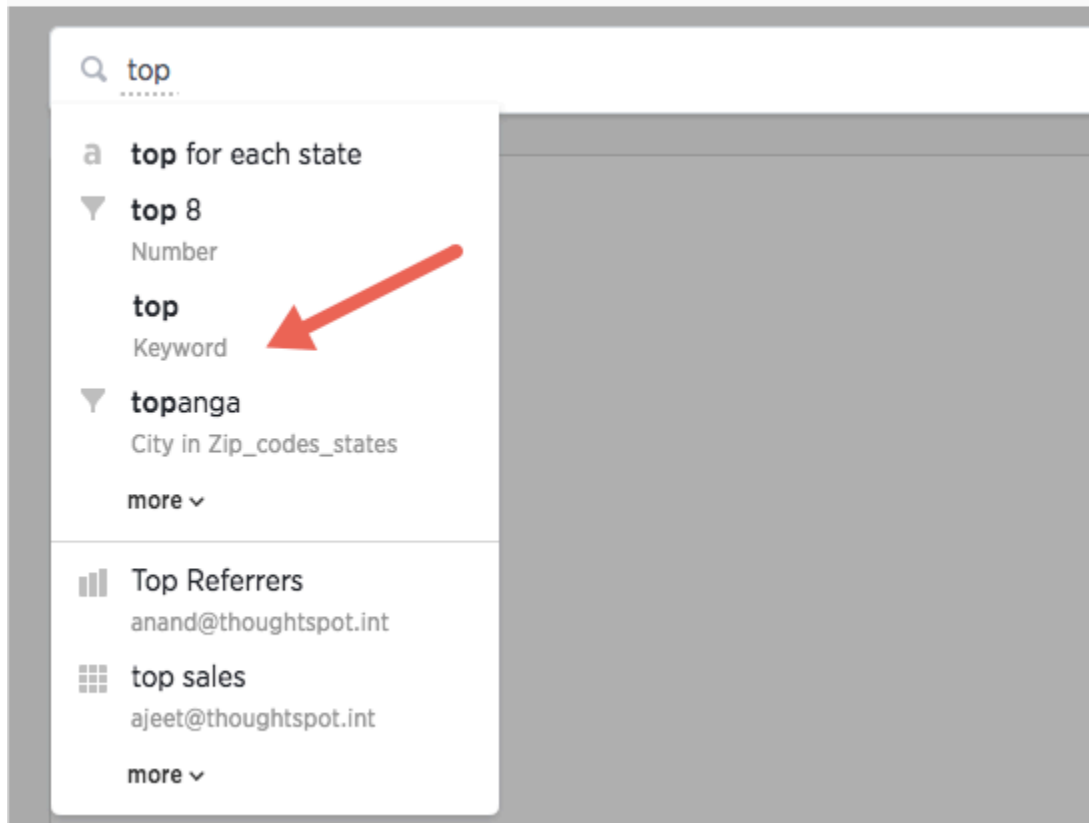
Now the `tscli nas mount-cifs` command prompts for a password if you do not supply one.

### *Sankey Chart*

A new type of chart, the Sankey diagram, is now available. This chart displays a flow diagram that illustrates transfers of flows within a system.

### *New improved icons*

In this release, the search ? (question mark) icon was replaced with an i (letter "i"). The team has redesigned the caret-down and x (clear-all) icon. Finally, the wand icon is removed from keywords in the **Search** dropdown. Instead, the identifier `keyword` appears beneath each keyword.

**Description**

### *What am I looking at? Redesign!*

The **What am I looking at** information panel is improved in this release with an **I** rather than the previous **?** (question mark). The new design is intended to clarify the information in the panel and make it easier to use.

### *Configurable maximum job time*

You can set the maximum amount of time in milliseconds a task is allowed to stay in "In progress" state. If the task exceeds this time, it is aborted.

### *Administer row level security privilege*

Users with the **Can Administer RLS** privilege can bypass row-level security (RLS) rules. By default this privilege is enabled. You can configure your installation to disable it.

## Description

ThoughtSpot introduced a feature in **3.4** which implicitly converted boolean columns to filters when a user searched. For example, `product isRed = true` became `product isRed`. This behavior meant that the `product` column was not added to the table or chart result.

Users were confused by this implicit behavior. This feature was reverted **4.3**. Now, boolean columns are treated like any other columns in the query. Meaning they remain `product isRed = true` and are not converted. The column also appears in the search results.

When you upgrade a 3.4 through 4.2 installation to 4.3 or newer, tables or charts created in the 3.4 through 4.2, that contained these booleans are upgraded using `product isRed = true` so that users are not surprised by additional columns suddenly appearing in their saved answers.

---

### *Improvements in RAMDisk recovery*

This release includes improvements in the speed and performance of RAMDisk recovery.

---

### *Fan trap situations are supported for new installations*

ThoughtSpot now supports fan traps in new installations. This feature is enabled by default in new installations.

For existing customers upgrading to 4.4 this feature is disabled. You must work with ThoughtSpot Support and Customer support to identify answers and pinboards that are impacted should this feature be enabled. Your ThoughtSpot customer service representative or ThoughtSpot Support can help you to identify, mitigate, and resolve any impacted data. Once these impacts are resolved, they can enable the feature for you.

---

### *Improvements to concat function*

The `concat` function now allows users to concatenate multiple strings rather than single strings.

---

### *Primary key drops check dependencies*

Dropping a primary key can impact downstream objects such as worksheets and pinboards. When dropping a primary key, TQL now checks for dependencies and issues a warning. Users must override with the `--allow_unsafe` flag.

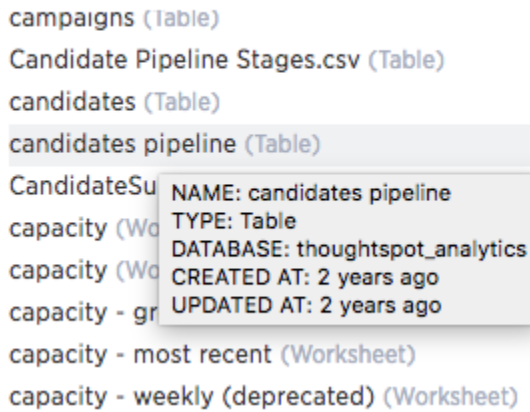
## Description

### *Tooltip functionality in the Schema Visualizer*

A tooltip function was introduced into the **Table** list for the **Schema Visualizer**. This tooltip shows information about the underlying table.

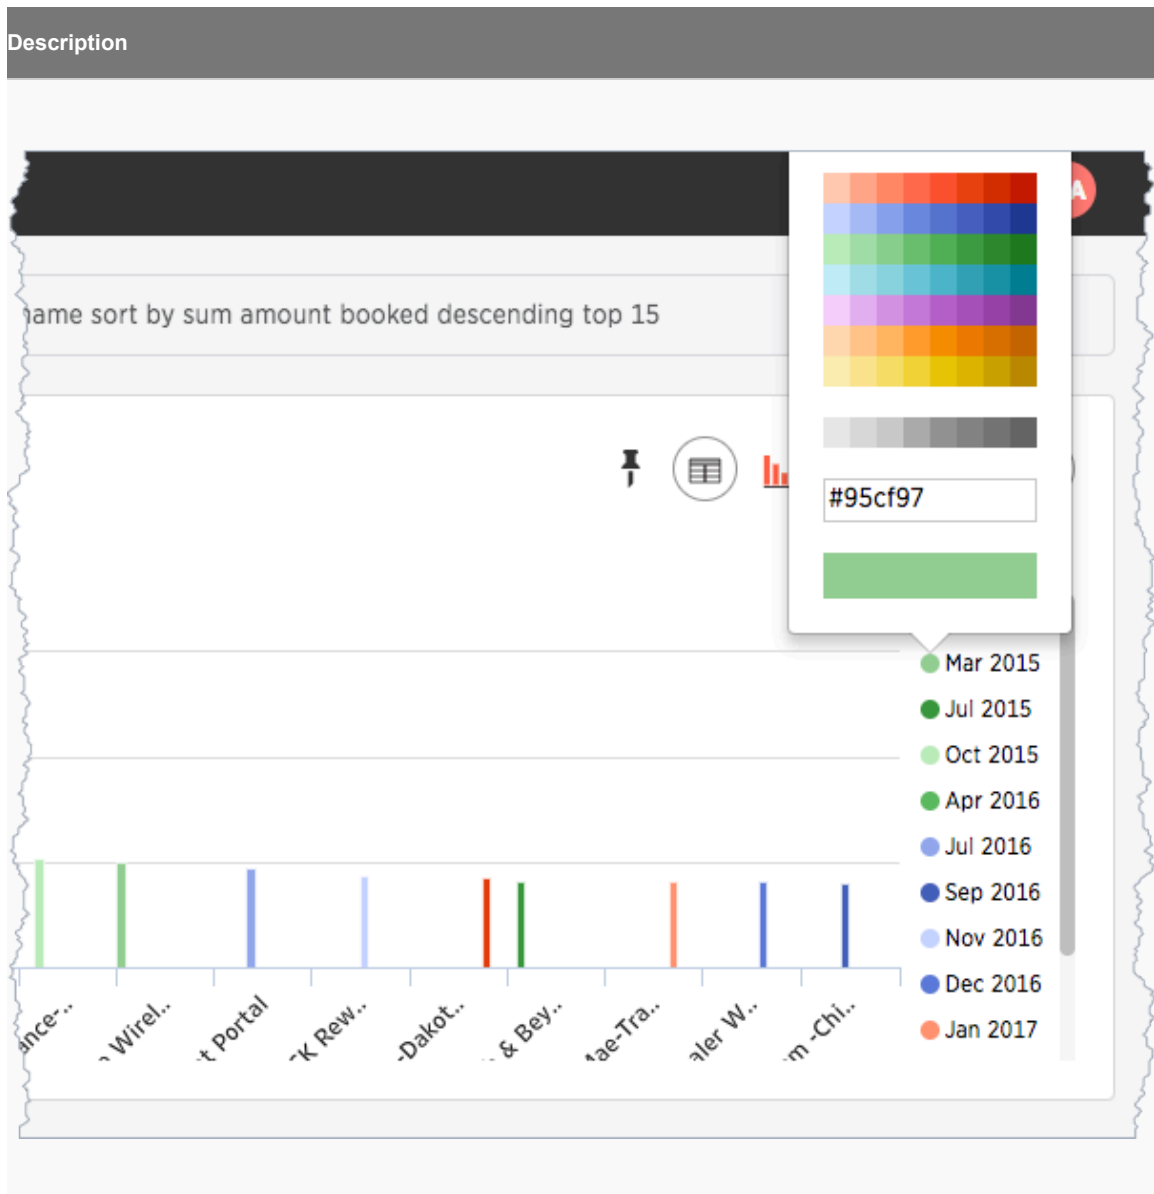
#### **Table**

campaigns (Table)  
Candidate Pipeline Stages.csv (Table)  
candidates (Table)  
candidates pipeline (Table)  
CandidateSu  
capacity (Wo  
capacity (Wo  
capacity - gr  
capacity - most recent (Worksheet)  
capacity - weekly (deprecated) (Worksheet)

A screenshot of a table list in the Schema Visualizer. The list contains several items, including 'candidates pipeline (Table)'. A tooltip is displayed over this item, showing the following information: NAME: candidates pipeline, TYPE: Table, DATABASE: thoughtspot\_analytics, CREATED AT: 2 years ago, and UPDATED AT: 2 years ago. The tooltip is a light gray box with a thin border and a drop shadow.

### *Color picker improvements*

This release includes new vibrant chart colors. It also includes the ability to add your own HEX values.



*tspublic session/login and session/logout APIs*

This release includes two new public APIs. You can use these APIs to log users in and out of a ThoughtSpot session. See the documentation for [details about the calls](#).

*Reduce search suggestion tokens*

This release improves the search experience by reducing the number of tokens in search results. Users found the number of suggestions confusing. Instead, the search greatly reduces the number of suggestions.

## Description

### Query processing improvements

This release includes improvements to the underlying processing of complex schema, object, and formula variations. Users should see improved ability to work with:

- Schema various such as chasm, fan, nested chasm, and chain of chasm traps,
- Object variations such as aggregated worksheet, join of aggregated worksheet and base tables.
- Formula variation such as aggregate and non-aggregate functions, windowing functions, and pinned measures.

These changes include optimization for generating queries that involve row level security(RLS).

### Filter by creator

Users have the ability to filter pinboards and answers by creator.

### Improvement in pinboard layout of small table

Tables with fewer rows appeared distorted with whitespace when pinned in a pinboard. This release includes improvements for display of tables in pinboards. See the documentation for more information about [editing a pinboard](#).

### Formula performance improvements

This release includes improvements in backend indexing formats that result in formula performance improvements.

## Description

### *Ability to define a join between a worksheet and a table*

Users can join a worksheet to one or more tables. The direction of the join can be from the worksheet to the table or the reverse. When you join a worksheet to a table, you bring any tables connected to the joined table into a search.

This feature has limitations. It is not possible to join multiple tables in the worksheet and a single table in a single relationship. Nor is it possible to join between a formula column in a worksheet and another table.

### *Getting a cluster configuration*

When using `tscli cluster get-config` to get cluster information, the system can run into situations where it cannot make a connection across the network. When this happens, the command returns an error but continues to work.

### *ODBC extract/transform/load improvements*

By default, ThoughtSpot takes a permissive approach to datatype compatibility between source and target in ODBC. As much as possible ThoughtSpot automatically converts incoming "compatible" input data to the desired ThoughtSpot target datatype. Flags were added that allow you to tune this behavior and make it stricter. At installation time, your customer support engineer can assist you in configuring a stricter load behavior.

### *Users with data access can drill down*

If you own the pinboard you can drill down to the data beneath. New in this release, users you share the pinboard with can also drill down provide they *also* have access to the data on which the board was based.

### *Description included when adding columns to worksheets*

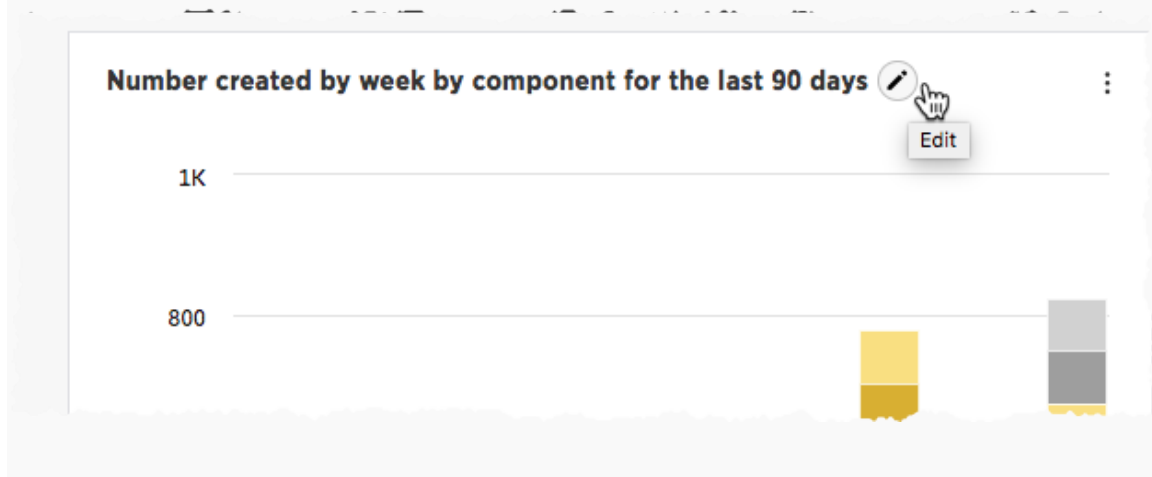
When adding a column from an existing data table or worksheet, the **Description** of the column is also copied into the new worksheet.



## Description

### Ability to change a pinboard headline

Users can edit the pinboard headline. To edit a headline, activate the headline edit button at the corner of a headline and click it.



### API to transfer ownership of objects

The new `POST /tspublic/v1/user/transfer/ownership` API was added. It allows the transfer ownership of all objects from one user to another. You must transfer all of a user's objects. This method takes two query parameters of type String.

- **fromUserName** Name of a user from whom the ownership needs to be transferred.
- **toUserName** Name of a user to whom the ownership needs to be transferred.

It returns a `200` status if ownership of all objects are successfully transferred. It returns a status of `400` if either the `fromName` and `toName` is invalid or for given username there is no user present.

## Fixed Bugs

### Description

Aggregation over `group_max` returned a red bar error, This was resolved, users can now aggregate over a `group_max` function.

A problem with editing formulas was resolved where users were repeatedly and unnecessarily presented with disam-

## Description

biguation options.

---

An issue was resolved where a bar on the chart label did not match the underlying data. The data was correct, the label was not.

---

An issue was corrected where, if one snapshot deletion failed in any fashion, any subsequent attempts to delete a snapshot failed.

---

Users were unable to edit a KPI formula if the KPI formula was pinned to a pinboard.

---

Users were able to **Select All** for 1000+ items in a pinboard filter. This action was available but not supported and caused users to believe pinboard filters were not working. Now, the option to **Select All** no longer functions when there are 1000+ items in a filter.

---

Users were unable to restore a database backup without first renaming the backup. Renaming is no longer required.

---

Creating a formula on `cumulative_sum` function mistakenly caused an error to appear. This action no longer causes an error.

---

Some installations saw worksheet performance degrade for worksheets with a large number of columns. Users can now set a configuration option to avoid these problems.

---

Previously setting both the browser and the ThoughtSpot profile value failed to display numbers and date formats in the proper locale. This problem was corrected. Setting the ThoughtSpot profile to the proper locale results in the appropriate display of date and number formats. It also causes translated strings to appear in the interface where they exist.

---

Active directory (AD) synchronization was not working. Moreover, if AD security group had no members in it, then the synchronization did not recognize the group at all. These problems were corrected, empty groups are recognized and users moved between groups now are properly synchronized by ThoughtSpot.

---

A problem was resolved where the round function returned a negative zero when it should have returned a 0 (zero).

---

## Description

Visualizations that relied on chasm trap worksheet were not immediately updated when a formula was changed. The workaround was to manually update the visualization. This was corrected. The system now updates the visualization automatically when an underlying formula is changed.

---

On a worksheet with a chasm trap, join between a fact and dimension did not work properly when grouping by a measure on the dimension table. This was fixed.

---

Using `safe_divide` and `sum` did not work with formula on formula. This problem was fixed.

---

Data connected schedules did not adjust for daylight savings time. Now, the schedules adjust as expected.

---

The ThoughtSpot tomcat instance went into crash loop when a generic relationship includes `IS NULL` operator. Now, this relationship no longer causes this problem.

---

An ODBC connection between Alteryx and ThoughtSpot was failing. This problem was solved by updating the ODBC drivers to the latest versions.

---

An issue was fixed where an unnecessary right-outer join was being applied to row-level security RLS-introduced joins. This situation led to impacts on query performance.

---

A problem was resolved with multiple sequential joins. Previously, the proper join path was not used when the fields are pulled from tables that were far apart unless fields from the intermediate tables were also used. Now, ThoughtSpot is able to handle multiple sequential joins.

---

Resolved an issue with a chasm trap that used `unique count` on shared dimension. The chasm trap was generating an incorrect query. Now, `unique count` is applied correctly in the final query.

---

Previously, a user user that created some row-level security (RLS) rules was associated with the rules such that deleting the user also deleted the rules. This is no longer the case. If a user creates a set of RLS rules and that user is subsequently deleted, the rules persist in the system.

---

An issue was resolved where filters could only be edited in the search bar.

---

## Description

Intermittent, ambiguous red bar errors related to chasm trap situations were resolved.

The help for row-level security rules was complicated and hard to parse. This text was improved and moved into the Query Visualizer.