



Release Notes

Version 3.5.7
February 2017

Contents

What's new in 3.5?	3
Scheduled pinboards	3
REST API pagination	5
What's changed in 3.5?	6
Clients and API files are separated from releases	6
Documentation changes	6
What's fixed in 3.5.7?	6
FIXED: IOException supportability error	6
FIXED: First time loading a pinboard is slow	6
FIXED: Data Connect - Data load fails due to invalid SSL cert error	6
FIXED: Security vulnerability with Tomcat version	7
FIXED: Day formulas are not working	7
FIXED: RLS query performance issues	7
FIXED: Logging in using single sign-on is broken	7
What's fixed in 3.5.6?	7
FIXED: Slow query time on worksheets with RLS	7
FIXED: Formulas in chasm trap don't work properly	7
FIXED: Cannot add relationships to old tables	7
FIXED: Not able to add RLS due to blank join mapping path	8
What's fixed in 3.5.5?	8
Support for LDAP non-anonymous bind	8
FIXED: Day of the week filter sorts incorrectly	8
FIXED: Answer with group_sum formula does not open	8
FIXED: Extra text is added to the end of a search string	8
FIXED: Visualizations cannot be recovered if underlying objects are deleted	8
FIXED: Existing formulas are removed on addition/removal of columns	9
FIXED: Unable to query on columns of two tables in a worksheet	9
FIXED: Worksheet is fetched partially due to root/leaf relationship	9
FIXED: "join path conflict" when joining tables into worksheets	9
FIXED: Add node fails due to failure of offline install	9
FIXED: Non-determinism causes longer query times	9
What's fixed in 3.5.4?	10
Bulk User/group synchronization APIs	10

What's fixed in 3.5.3?	10
Ability to dynamically embed data	10
Ability to search header lists by name and display name	10
Propagate errors to parent app in embedding	10
FIXED: Download as Excel file preserves "Week" keyword	10
FIXED: Cumulative sum returns incorrect results when saved as an aggregated worksheet	11
FIXED: Failure to download a pinboard as a PDF	11
FIXED: Make pinboard title dropdown look consistent	11
What's fixed in 3.5.2?	11
Embedded content creation	11
Formula indexing	11
What's fixed in 3.5.1?	12
Download a pinboard as a PDF	12
FIXED: Data load timestamp is unavailable on column tooltip	12
What's fixed in 3.5?	12
FIXED: Data filter for time is not supported	12
FIXED: Version number in the help center shouldn't come from browser cache	13
Known issues	13
Applying filter on {blank} value results in no data	13
Aggregation on top of an aggregated formula error	13
Proper error message for tsload errors	13
tscli monitoring command should check validity of email address	13
Case sensitivity not preserved when importing URLs	13
KPI line extends outside of the chart canvas	13
Date format issues in pivot table	14
Supported upgrade paths	14

What's new in 3.5?

Scheduled pinboards

You can now get your pinboards emailed to you on a regular basis to use for offline analysis. This introduces a new format for you to consume and share pinboards with others, including those who don't have a ThoughtSpot account. Scheduled pinboards should help with preparing for recurrent meetings, when reviewing the same pinboard is necessary.

You can view all of the schedules set for a pinboard under Actions.

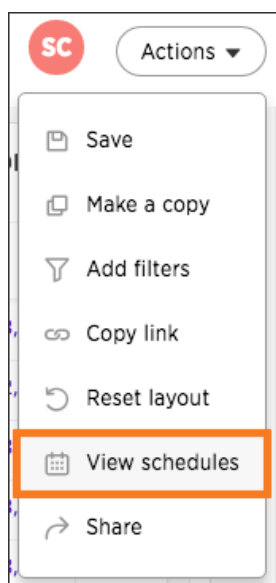


Figure 1: View pinboard schedules

Or you can view all scheduled jobs on the Jobs Management page under Admin.

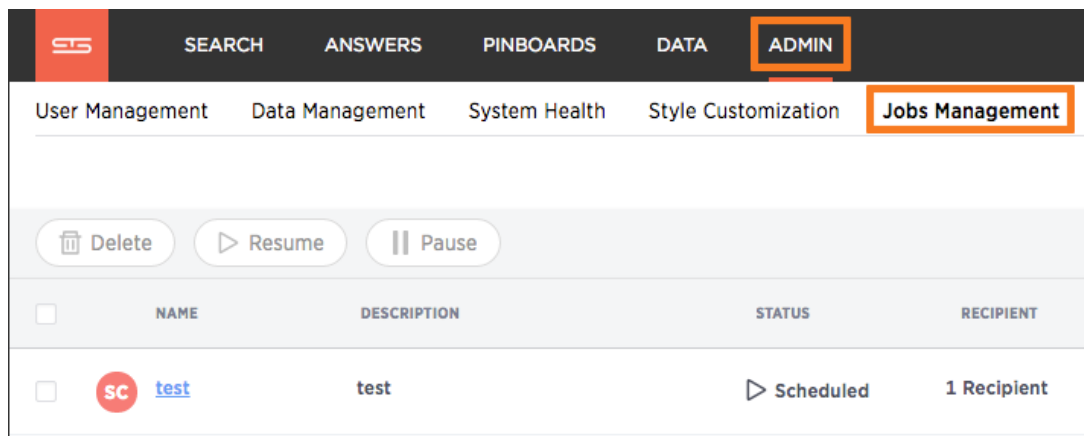


Figure 2: View Jobs Management page

Scheduled pinboards can be set to repeat every n minutes, hourly, daily, weekly, or monthly. For some of these, you can also choose specific times of the day or days of the week.

Schedule

Repeats Weekly

Run the task every week on

Sunday Monday Tuesday Wednesday Thursday

Friday Saturday

at 12 : 00

Figure 3: Set pinboard schedule

The pinboard data is sent as either a CSV or PDF file, with one pinboard scheduled per job. CSV files provide all data for tables, with one attachment provided per table. Use CSV files to perform further analysis offline. PDF files show all visualizations in the pinboard. Each chart takes up a whole page in the file, and the first 100 rows of a table are included. Use PDF files to skim the data.

Server time zone America/Los_Angeles

Name Weekly Pinboard Report

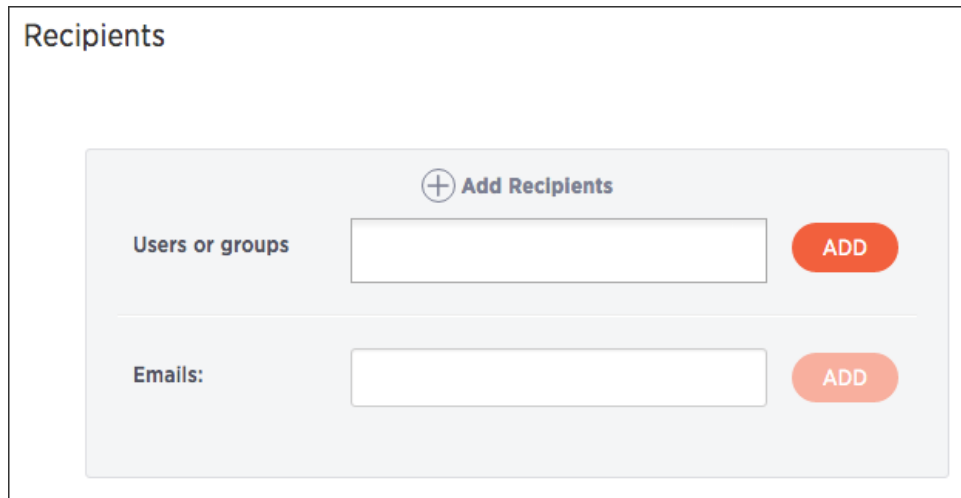
Description Weekly pinboard report schedule for sales.

Type CSV PDF

CSV files will be sent only for the tables in this pinboard.

Figure 4: Set pinboard file type

You can add ThoughtSpot users as recipients, and if you have whitelist domains set, you can even add email addresses of users who aren't on ThoughtSpot. You can add up to 1000 recipients.



The image shows a web interface for adding recipients. At the top, the word "Recipients" is displayed. Below it is a light gray box containing a "+ Add Recipients" button. Underneath this box, there are two rows of input fields. The first row is labeled "Users or groups" and has a white text input field followed by a red "ADD" button. The second row is labeled "Emails:" and has a white text input field followed by a red "ADD" button.

Figure 5: Set pinboard recipients

Some additional information about scheduled pinboards includes the following:

- Administrators and users with can schedule pinboard privilege can schedule and manage pinboard schedules.
- The scheduled pinboards feature respects row level security rules. This means if the recipient of a pinboard is a user in ThoughtSpot, then they can only see data based on their own access to the pinboard. If the recipient is from outside of the cluster, then they will have access to the dataset of the pinboard based on the sender's permissions.
- You can have up to 50 scheduled jobs on a cluster.
- The size of each scheduled pinboard email can be as large as 25 MB, which matches most email services size limitations.

You can find more details on how to create and manage scheduled pinboards in the [ThoughtSpot Administrator Guide](#).

REST API pagination

When pulling data out of ThoughtSpot and into your own portal, you can paginate the data to decrease load time. The data that you call from the REST API would be paginated so that your data tables are populated quickly. And every time you go to the next page of data on the table, a new REST call would be made.

You can find more details in the [ThoughtSpot Application Integration Guide](#).

What's changed in 3.5?

Clients and API files are separated from releases

The ThoughtSpot clients (ODBC and JDBC drivers) and API files have been moved out of the release. This not only reduces the tarball size, but also allows ThoughtSpot to update the clients outside of the release cycle to fix bugs.

You can continue to find the clients and API files in the Downloads section of the Help Center.

Documentation changes

There is a new guide available in the ThoughtSpot documentation library:

- [ThoughtSpot on AWS Quick Start Guide](#): This new quick start guide includes information on how to launch ThoughtSpot on AWS.

Information about this topic has been removed from the [ThoughtSpot Administrator Guide](#) and placed in the appropriate new guide.

- **About restore operations**: This topic has been moved to the [ThoughtSpot Internal and POC Guide](#). Please contact ThoughtSpot Support if you would like to restore from a snapshot or backup.

What's fixed in 3.5.7?

FIXED: IOException supportability error

ISSUE ID: SCAL-13435

Description: There was an IOException when the system attempted to append to the event log, affecting general supportability.

FIXED: First time loading a pinboard is slow

ISSUE ID: SCAL-17619

Description: Loading certain pinboards for the first time would occasionally take several seconds to complete, while subsequent loads were near-instantaneous. This was due to a query caching issue.

FIXED: Data Connect - Data load fails due to invalid SSL cert error

ISSUE ID: SCAL-17722

Description: When trying to load data using Data Connect, the data load would fail due to an invalid SSL cert error, even after the SSL validation check was bypassed.

FIXED: Security vulnerability with Tomcat version

ISSUE ID: SCAL-18018

Description: We have updated Apache Tomcat to version 8.5.11.

FIXED: Day formulas are not working

ISSUE ID: SCAL-18165

Description: The day and hour_of_day formulas would not show up in a worksheet after creating them. You could only see them in the left panel when searching, but even then, they would return blank results. This has been fixed.

FIXED: RLS query performance issues

ISSUE ID: SCAL-18212

Description: RLS rules that used strcat resulted in slow query times. The workaround had you rewrite the RLS rules so that they did not incorporate strcat.

FIXED: Logging in using single sign-on is broken

ISSUE ID: SCAL-18391

Description: The sign in link wouldn't show up on the login page even after performing SAML configuration. This display issue has been fixed.

What's fixed in 3.5.6?

FIXED: Slow query time on worksheets with RLS

ISSUE ID: SCAL-17955

Description: There was a significant delay when querying on a column from a fact and dimension table, where the tables were linked in a worksheet with RLS and non-progressive joins. The performance has been improved.

FIXED: Formulas in chasm trap don't work properly

ISSUE ID: SCAL-18077

Description: Formulas that were part of a data model with two fact tables joined by a chasm bridge did not work properly. The formulas in the worksheet contained counts on the main fact tables. When these formula fields were dragged into the search bar, a wrong answer would be returned. This was because the subqueries were not being properly joined. However, entering the count fields directly into the search bar would return the correct answer.

FIXED: Cannot add relationships to old tables

ISSUE ID: SCAL-18194

Description: It was not possible to add relationships to system tables created before ThoughtSpot version 2.3 since they were missing proper schema and database configurations. These tables had to be manually fixed.

FIXED: Not able to add RLS due to blank join mapping path

ISSUE ID: SCAL-18236

Description: You were not able to add RLS on a column of two tables due to a blank join mapping path. This issue persisted even when a join path existed between the tables. This bug has been fixed.

What's fixed in 3.5.5?

Support for LDAP non-anonymous bind

ISSUE ID: SCAL-17796

Description: ThoughtSpot now supports non-anonymous LDAP integration. Non-anonymous LDAP binding is more rigorous than anonymous authentication, but it should help you track what your users are querying and keep a log trace for auditing purposes.

FIXED: Day of the week filter sorts incorrectly

ISSUE ID: SCAL-7284

Description: Filtering by the day of the week did not sort the days based on calendar day order. This has been fixed.

FIXED: Answer with group_sum formula does not open

ISSUE ID: SCAL-15063

Description: An answer with a group_sum formula would return an error and not open. This has been fixed.

FIXED: Extra text is added to the end of a search string

ISSUE ID: SCAL-17730

Description: Clicking a search suggestion sometimes added a duplicate search token to the end of your search string in the search bar. This extra text had to be manually removed before you could resume your search.

FIXED: Visualizations cannot be recovered if underlying objects are deleted

ISSUE ID: SCAL-17756

Description: There was no way to recover dependent answers and pinboards if an underlying data source, such as a column, was deleted by mistake. The lack of a recovery mechanism meant

you lost all of your work. However, now you can recover the affected visualizations by editing them. This allows you to change the query and remove sources that no longer exist.

FIXED: Existing formulas are removed on addition/removal of columns

ISSUE ID: SCAL-17757

Description: Adding or removing a column or adding another formula to your answer would remove existing formulas. This has been fixed.

FIXED: Unable to query on columns of two tables in a worksheet

ISSUE ID: SCAL-17959

Description: Querying on columns of two tables in a worksheet returned an error that the data was still being indexed. This was due to a missing request in the system which allowed this error to show up indefinitely. The tables had been sharded, and the relationships and constraints were modified to take the sharding into account.

FIXED: Worksheet is fetched partially due to root/leaf relationship

ISSUE ID: SCAL-17962

Description: Creating a worksheet using columns from only the root and leaf tables of three tables with consecutive relationships would cause errors. The worksheet would be marked as incomplete when you tried to retrieve it.

FIXED: “join path conflict” when joining tables into worksheets

ISSUE ID: SCAL-17973

Description: There was a bug that would add a logical table with an empty join path. If two tables with empty join paths were added, then the second table would run into a join path conflict.

FIXED: Add node fails due to failure of offline install

ISSUE ID: SCAL-18023

Description: A failed offline install also led to a failed node addition. The workaround had you edit a setup.py function.

FIXED: Non-determinism causes longer query times

ISSUE ID: SCAL-18078

Description: It would sometimes take 10 to 20 seconds to receive answers to questions over small amounts of data. This was due to a lack of caching and non-determinism in the order in which joins were being generated.

What's fixed in 3.5.4?

Bulk User/group synchronization APIs

ISSUE ID: SCAL-16302

Description: You can use public APIs to customize authentication and authorization for large numbers of users and groups. These actions include being able to bulk create/delete users and groups, assign/remove users to groups, and search for users based on user ID or group membership.

What's fixed in 3.5.3?

Ability to dynamically embed data

ISSUE ID: SCAL-16198

Description: You can make a call to get the list of content items, answers and pinboards, then implement selectors to let users select the content to view.

We have made the metadata list API endpoint public, and have added a new public endpoint to fetch headers of visualizations in a given pinboard. The information returned consists of the title, GUID, size, and any other simple header information of the visualization.

Ability to search header lists by name and display name

ISSUE ID: SCAL-17461

Description: You can now search header lists by both name and display name. This should be especially useful when searching for users or groups on the User Management page.

Propagate errors to parent app in embedding

ISSUE ID: SCAL-17546

Description: In the embed use case, you can now intercept ThoughtSpot's error messages and communicate them to your users in a way that makes sense in the context of your parent app. This is done by suppressing error messages in the UI, and passing their details along to your parent app.

FIXED: Download as Excel file preserves "Week" keyword

ISSUE ID: SCAL-14529

Description: Downloading your table as an Excel file preserved the "Week" keyword in the date, "MM/DD/YYYY Week" column. This made it difficult to perform analysis on the data. Now, all of

the date/date-time fields present in a downloaded Excel file have date cell type, and the keyword is no longer preserved.

FIXED: Cumulative sum returns incorrect results when saved as an aggregated worksheet

ISSUE ID: SCAL-17385

Description: Cumulative sum formulas returned incorrect results when they were saved as part of an aggregated worksheet. This bug was only reproducible when the date column was saved without any bucketing.

FIXED: Failure to download a pinboard as a PDF

ISSUE ID: SCAL-17501

Description: Downloading a pinboard as a PDF would sometimes timeout and fail. This was usually due to the feature downloading all of the rows of a table visualization, resulting in thousands of pages, instead of only the first 100.

FIXED: Make pinboard title dropdown look consistent

ISSUE ID: SCAL-17547

Description: The pinboard dropdown widget looked different on the pinboard page when compared to the homepage. This has been fixed so the look is consistent and clear that it is a dropdown widget.

What's fixed in 3.5.2?

Embedded content creation

ISSUE ID: SCAL-16199

Description: This full embed feature enhances our existing external sharing functionality. You can create content in an embedded environment, which includes being able to ask questions, create answers and pinboards, and analyze data.

Formula indexing

ISSUE ID: SCAL-17404

Description: In 3.4, indexing data values for formula columns was disabled by default. This was done to prevent performance issues caused by indexing queries for some of these formulas. In this patch, indexing is re-enabled for formulas that are unlikely to cause issues. These include formulas whose output space is restricted to hard-coded strings and those that do not require any joins.

What's fixed in 3.5.1?

Download a pinboard as a PDF

ISSUE ID: SCAL-16834

Description: You can download a pinboard as a PDF, without having to download each visualization separately. It works just as you would download an answer. To do so, while viewing your pinboard, click **Actions** and select **Download as PDF**.

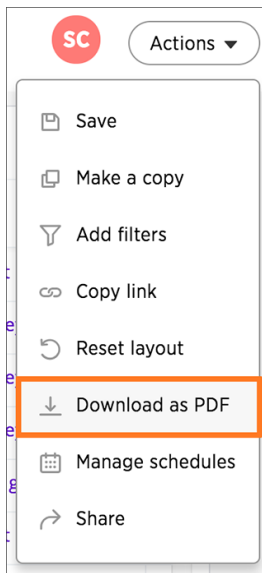


Figure 6: Download as PDF dropdown option

A one time PDF file will be generated with all of the same specifications as if you received it in a scheduled pinboard email.

FIXED: Data load timestamp is unavailable on column tooltip

ISSUE ID: SCAL-16944

Description: In the chart column name tooltip, information on when the data source was last updated was missing.

What's fixed in 3.5?

FIXED: Data filter for time is not supported

ISSUE ID: SCAL-8579

Description: Data filtering in a table column did not allow you to properly filter time values.

FIXED: Version number in the help center shouldn't come from browser cache

ISSUE ID: SCAL-13972

Description: The version number displayed in the bottom right of the help center would sometimes be incorrect. This was because the information was taken from browser cache. You had to clear browser cache to see the correct number. Now, ThoughtSpot always shows the correct release number.

Known issues

Applying filter on {blank} value results in no data

ISSUE ID: SCAL-7731

Description: Applying a filter on {blank} values results in an answer with no data.

Aggregation on top of an aggregated formula error

ISSUE ID: SCAL-12193

Description: An error message is returned when you perform an aggregation, such as sum, on top of an aggregated formula.

Proper error message for tsload errors

ISSUE ID: SCAL-13149

Description: Error messages when loading data using tsload are not accurate.

tscli monitoring command should check validity of email address

ISSUE ID: SCAL-14471

Description: `tscli monitoring set-config --email` does not check the validity of email addresses. Instead, the command takes whatever value is enter.

Case sensitivity not preserved when importing URLs

ISSUE ID: SCAL-14588

Description: Case sensitivity is not preserved when importing URLs to the data view. This issue is important since some URLs require case sensitivity.

KPI line extends outside of the chart canvas

ISSUE ID: SCAL-14788

Description: The range of the y-axis is not taking the KPI value into consideration, so the line is beyond the scope of the chart canvas.

Date format issues in pivot table

ISSUE ID: SCAL-17160

Description: Date formats that are displayed correctly in the data view are corrupted in pivot tables.

Supported upgrade paths

If you are running a release that begins with one of the following ThoughtSpot version numbers, you can upgrade to this release directly:

- 3.4
- 3.5

If you are running a different version, you need to do a multiple pass upgrade. First, upgrade to one of the above versions, and then to this version.