

Data Pipeline Platform Comparison: Fivetran, Matillion, Stitch

GigaOm evaluates three major data pipeline platforms



Learn how Fivetran data integration powers
business intelligence at fivetran.com →

GIGAOM

BENCHMARK/FIELD SERVICE

Data Pipeline Platform Comparison **v1.0**

Fivetran, Matillion, Stitch

WILLIAM MCKNIGHT AND JAKE DOLEZAL



CREDIT: PETE LINFORTH FROM PIXABAY

Data Pipeline Platform Comparison

Fivetran, Matillion, Stitch

TABLE OF CONTENTS

- 1 Summary
- 2 Data Pipeline Platforms
- 3 Field Test Setup
- 4 Field Test Results
- 5 Conclusion
- 6 Disclaimer
- 7 About Fivetran
- 8 About William McKnight
- 9 About Jake Dolezal
- 10 About GigaOm
- 11 Copyright

1. Summary

Data is the currency of digital transformation. Having available data that is understood, organized, and believable strengthens all major corporate initiatives. However, maintaining this basic resource is a growing challenge for most organizations because sources and volumes of interesting data are expanding rapidly. The cloud and the proliferation of SaaS companies has contributed to the data explosion. While the possibilities of the cloud and its many applications can quickly grow the capabilities of an organization, the data spread it creates can lead to problems such as decentralized data leading to inaccurate findings, or wasted time spent rebuilding pipelines instead of driving results.

Without robust automation, an organization's data movement needs can quickly outpace the ability of a data engineering staff to meet that need.

Given growing workloads and a lack of data engineering resources, automation and ease of use are fundamentally important. Data pipelines are one aspect of the modern data stack that can be automated to solve for this growing challenge.

In this report, we compare the three major data pipeline platforms: Matillion, Stitch, and Fivetran; and run them through a series of selected tests that highlight their degree of automation, ease of setup, and documentation. We evaluated aspects that include the time and effort required to set up a source-destination connection, the degree of automation throughout the process, and the quality of documentation to support the effort. These areas address the three major "humps of work" we have encountered in our field work with data pipelines.

Of the three offerings, Fivetran had the shortest and easiest setup. Matillion Data Loader produced the longest setup, with the most steps. Matillion also had some steps that were poorly documented. Stitch ranked between Fivetran and Matillion Data Loader in our assessment, but it had the longest-running individual task (selecting which Salesforce entities to sync).

Fivetran handled the data source changes with full automation, while Matillion Data Loader presented the biggest automation challenge. Not only did the new data/altered columns not appear automatically in Matillion, but the pipeline had to be rebuilt. Stitch likewise required manual intervention to work with new data/altered columns.

Fivetran had the most thorough documentation across all the items we measured. Stitch also had good documentation, with only a few items either omitted or left short. Matillion Data Loader's documentation describing the data of the source data connector for Salesforce was nearly completely missing.

Another observation: We found the level of loading and updating activity in Snowflake caused by the Matillion solution to be excessive compared to Stitch and Fivetran. Data pipelines like these are well worth exploring for any enterprise data integration effort, especially where your source and target are supported.

2. Data Pipeline Platforms

At first glance it may be difficult to distinguish among these products. They are all Extract, Load, Transform (ELT) solutions, and they are all cloud-based, cloud-native, with a visual authoring environment and a host of source and target connectors. In our field test, we provide more color to these products in specific and the data pipeline market in general.

Matillion

The Matillion products that work together for the data pipeline solution are Matillion ELT and Matillion Data Loader. It comes with a visual authoring environment and fully leverages ELT by keenly utilizing the newly-minted power of the cloud database platforms to which it serves data. Matillion leverages the massively parallel processing (MPP)-based bulk load capabilities of the target by decoupling the load from the transform steps. In this report, we tested only Matillion Data Loader, because the scope of the report tested as-is bulk data movements. Matillion has connectors for over 70 well-cultivated data sources and supports Redshift, BigQuery, and Snowflake targets.

Stitch

Stitch, owned by Talend, is a data pipeline product with a high number of connectors (90-plus), with many of these being community supported. Like the other products in this field test, Stitch allows you to focus on data analysis instead of a prolonged build phase. Stitch connects to many Software as a Service (SaaS) applications and is both relatively easy to use and quick to provision compared to prior generations of data integration.

Fivetran

Fivetran takes an expansive view of quick provisioning and productivity. Automation and ease of use are prevalent in Fivetran across all functions. The solution boasts more than 150 connectors, and supports modern destinations such as Databricks, BigQuery, Snowflake, Redshift, Azure SQL Data Warehouse, and more. Its built-in schema change detection and propagation saves time and ensures the long-term accuracy of its low-maintenance pipelines while also supporting the automation of many non-native connections.

3. Field Test Setup

The field test was designed to assess the capabilities, features, ease-of-use, and documentation of the three data pipeline platforms. These are the four major items we find are critical to having success with data pipelines. Of course, testing of this nature is very challenging. We strove to eliminate as much subjectivity as possible from the test plan, methodology, and measurement. However, we concede that different test configurations can favor one vendor over another by the design of the test itself. Our testing demonstrates a narrow slice of potential configurations and scenarios.

GigaOm partnered with Fivetran, the sponsor of this report, to select competitive platforms that offer comparable features and capabilities to address organizations' data pipeline use cases. GigaOm selected the test scenario, methodology, and configuration of the environments.

We leave the issue of the fairness of the report for the reader to determine. We strongly encourage you, as the reader, to look past marketing messages and discern for yourself what is of value. We hope this report is informative and helpful in uncovering some of the challenges and nuances involved in platform selection.

The parameters used to replicate this benchmark are provided throughout this document. We have provided enough information in the report for anyone to reproduce this test. We encourage you to compile your own representative use case and test compatible configurations applicable to your requirements.

Test Scenario

For the data pipeline platforms, we selected a simple and straightforward, but very common, use case for our testing. The scenario involves using each of the three data pipeline tools to initially lift and continually feed fluidly changing source data from a popular cloud-based customer relationship management (CRM) software tool and load it into a popular cloud-based data warehouse platform. In this case, we chose Salesforce as the CRM data source, and Snowflake as the data warehouse destination.

Salesforce represents a broad spectrum of similar use cases because it has both a conventional relational database and an API access layer. It also is used widely as a fully-managed cloud offering. Thus, it was a good candidate to represent the real-world scenario we tested. In the same vein, Snowflake is a popular, fully-managed cloud, columnar data warehousing platform. Of course, not all organizations will have these same technologies in place, but Salesforce and Snowflake technologies are representative of a broad use case.

Test Environment

The field test consisted of Salesforce CRM: Enterprise Edition as a source system. For the test, we generated test data from 10,000 Accounts, 50,000 Contacts, and 50,000 Leads. The data was

randomly generated and loaded into the Salesforce system using its Bulk Data Import utility.

For the destination, we created three separate databases, three schemas, and three users—one for each platform—within the Snowflake environment.

We then signed up for free trials of each of the competitive platforms—Fivetran, Stitch, and Matillion Data Loader. Each platform offers a limited-time, but fully featured, offering of their product. At the time of this report, Matillion does not offer a free trial of Matillion ELT. However, our test did not include any data transformations, only bulk data movement as-is.

We then followed the documentation of each data pipeline tool to configure the environment, set up the source connection, set up the destination connection, and begin syncing data from the source to the destination. Again, we did not perform any transformations to the data.

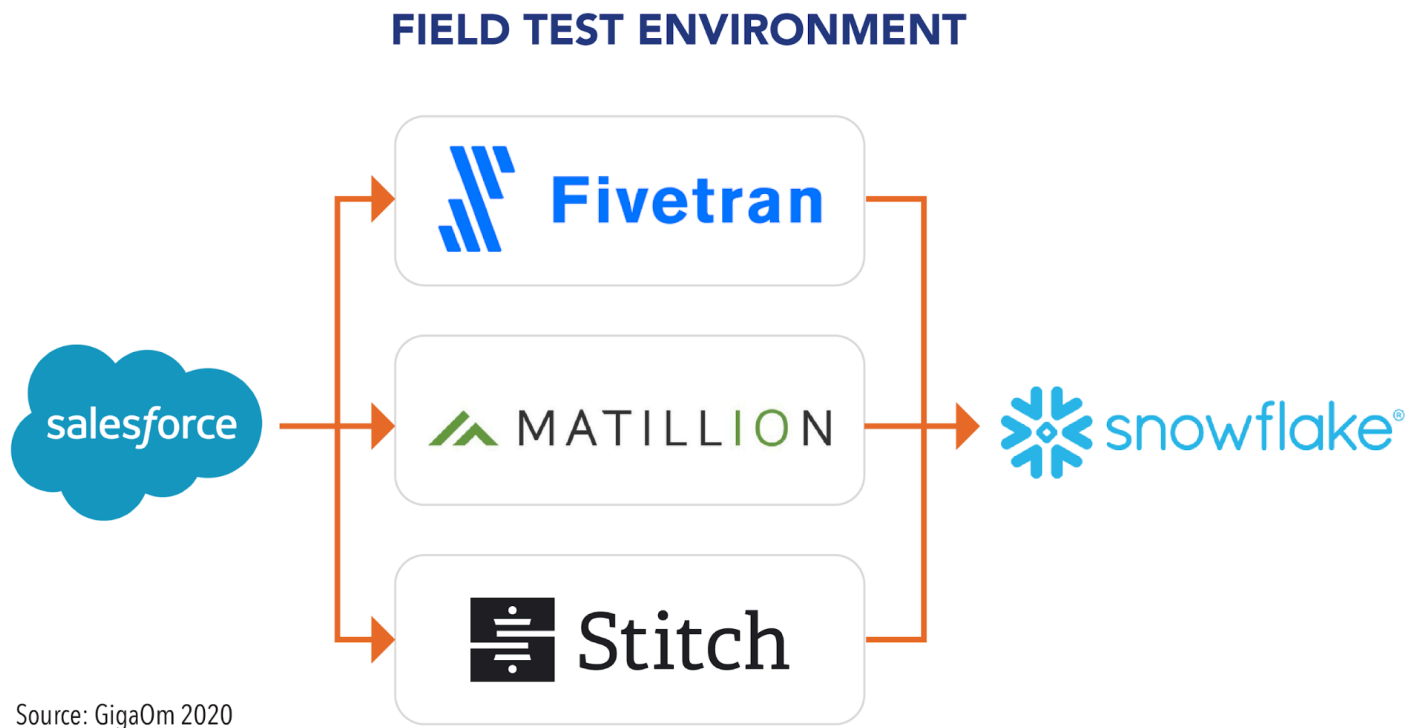


Figure 1. Field Test Environment

Test Methodology

The Field Test consisted of three separate tests:

1. Setup Test

2. Automation Test
3. Documentation Test

Test 1: Setup

The setup test measures the amount of effort or task complexity involved in setting up a source integration and a target destination from end to end and start to sync. It also measures the thoroughness in the documentation for each step in the process. The test has four measures:

1. Source Setup Effort – This measure is a combination of the number of tasks or steps and the length of time required to complete each task. The fewer number of tasks and the shorter the total amount of time required for all tasks, the higher the score.
2. Source Setup Documentation – Each task’s documentation on each platform’s public website is assessed into the following categories (from best to worst):
 1. Fully documented – All actions required to complete the step are accounted for and clearly described in the documents.
 2. Partially documented – All actions required to complete the step are accounted for, but some details are omitted.
 3. External reference – Actions are accounted for, but the user must visit another web page or site to complete the task.
 4. Missing – Some or all actions required are missing from the documents.
3. Destination Setup Effort – This measure is a combination of the number of tasks or steps and the length of time required to complete each task. The fewer number of tasks and the shorter the total amount of time for all tasks, the higher the score.
4. Destination Setup Documentation – Each task’s documentation on each platform’s public website is assessed into the following categories (from best to worst):
 - Fully documented – All actions required to complete the step are accounted for and clearly described in the documents.
 - Partially documented – All actions required to complete the step are accounted for, but some details are omitted.
 - External reference – Actions are accounted for, but the user must visit another web page or site to complete the task.
 - Missing – Some or all actions required are missing from the documents.

Test 2: Automation

The automation test measures the amount of effort required and the level of automation applied when changes are made to the source data model. The effort is assessed into the following categories (from best to worst):

- Fully automated – Changes are captured into new columns (and old data is preserved) without any action or effort on the part of the user.
- Manual – Changes must be identified and added to the sync integration by the user.
- Rebuild – The sync integration must be rebuilt from scratch.

The Automation test has four measures. They are:

1. Capture New Column – A new column is added to an existing entity in the source system.
2. Capture Renamed Column – An existing column is altered and renamed in an existing entity in the source system.
3. Capture Altered Column Data Type – An existing column's data type is altered in an existing entity in the source system.
 1. Retain Old Column After Alter/Rename – After the Rename/Alter in 2b and 2c, the destination column of the previous name/type is retained in its original form as syncing is performed on the new column.
4. Capture New Table – A new entity table is added in the source system.

Test 3: Documentation

The documentation test measures the comprehensiveness of the documentation covering both source and destination schema design, naming conventions, and load behavior. Products are graded across four categories (from best to worst):

- Fully documented – Schemas and entities are accounted for and clearly described in the documents.
- Partially documented – Schemas and entities are accounted for, but some details are omitted.
- External reference – Schemas and entities are accounted for, but the user must visit another web page or site to see them.
- Missing – Schemas and entities are missing from the documentation web site.

The Documentation test has four measures. They are:

1. Source/Destination Table Relations – An entity-relationship diagram showing the relationship between key tables in the Source and Destination tables
2. Entity Table List – A list of all entity tables that get synced by the connector
3. Data Source Usage – Specific information on the method and frequency with which the source system is accessed (such as API calls and any quotas that run the risk of being exceeded)
4. Connector Load Behavior – Information regarding changed naming conventions or other table import behavior

Test Scoring

The tests were scored on a scale from 1 to 5, where 5 is the highest (best), and 1 is the lowest score. All the scores were then averaged across all three tests, and results are shown within a rubric in the Field Test Results section of this document.

Setup Effort Scoring (Tests 1a and 1c)

To score tests 1a and 1c to measure setup effort, we measured the amount of time it took to complete each step when we followed each platform’s documentation. We also realize that as data practitioners, our professional experience with testing and data integration may introduce bias into measurements of the amount of time required. More novice users might take more time than we do, yet even more experienced or specialized data integration professionals may be able to complete the steps faster than we can. Thus, we developed a relative effort scale using T-shirt sizes—extra-small (XS), small (S), medium (M), large (L), and extra-large (XL). Depending on your experience, the actual time it takes may vary.

Thus, we assigned every task we completed a T-shirt size, and then we converted those to continuous integer values from 1-5, where an XS was a 1, and an XL task was a 5. We added those values for all the steps required to complete the overall setup and divided by the number of tasks to get an average, which we then rounded to the nearest integer. We then used the following scoring matrix of the total number of steps and the average task size of each step to arrive at the final score.

Table 1. Setup Effort Scoring Matrix

SETUP EFFORT SCORING MATRIX

AVERAGE TASK SIZE OF EACH STEP (ROUNDED)

TOTAL # OF STEPS	XS 1	S 2	M 3	L 4	XL 5
1-2	5	5	4	3	2
3	5	4	3	3	2
4	4	3	2	2	1
5	4	3	2	1	1
6+	3	2	1	1	1

Source: GigaOm 2020

Automation Tests (All of Test 2)

For all of Test 2, we used the following scores to assess the level of automation achieved by each platform according to the test methodology above.

- Fully automated – 5
- Manual – 3
- Rebuild – 1

Documentation Scoring (Tests 1b, 1d, and All of Test 3):

For Tests 1b and 1d, we gave each step/component the following scores and the results were totaled and averaged. For all of Test 3, we assigned one of these scores to each measure.

- Fully documented – 5
- Partially documented – 3
- External reference or missing – 1

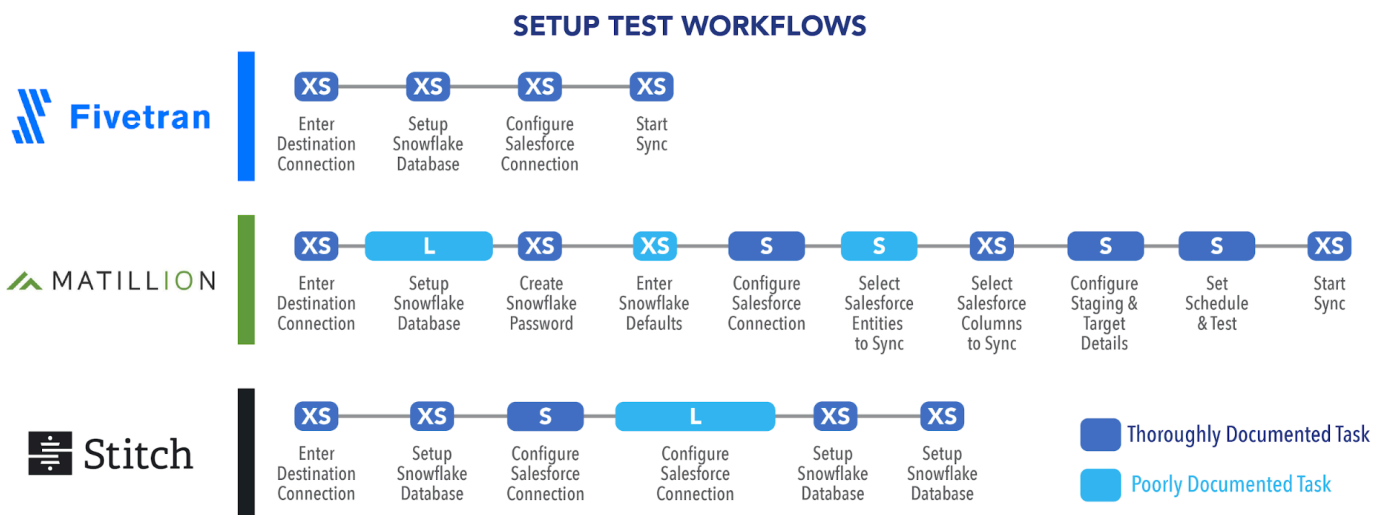
4. Field Test Results

The following section reveals and discusses the results we found when conducting the field test according to the methodology described above. Again, results may vary across different configurations and test scenarios. You are encouraged to compile your own representative use case and test compatible configurations applicable to your requirements. The following results are only a slice of potential outcomes.

Test 1: Setup Test Results

We set up each data pipeline platform according to its documentation. We recorded the time it took for each step and noted the level and completeness of the documentation in describing how to perform each step.

The following workflow diagram outlines what we uncovered.



Source: GigaOm 2020

Figure 2. Setup Test Workflows

The workflow diagram shows nodes that represent a single distinct step in the process. The size of each step, or the duration of time it took to complete, is represented by both the width of the node and the T-shirt size we assigned. Thus, the number of nodes and the length of the workflow demonstrate the ease of setup. Fewer nodes and shorter tasks are indicative of a simple setup. Also, dark-colored nodes represent tasks that were well documented—meaning everything we need to do to complete the step was clearly outlined in the documentation. A lighter-colored node was a poorly documented step, for which essential information needed to complete the step was either omitted or referenced to an external vendor’s website. Sending the user to Snowflake’s website for instructions on how to create a new schema is one example of the latter. These poorly documented steps also added to the time it took to complete the step because we (the user) had to decipher for ourselves how to fill in the

blanks.

As you can see, Fivetran has the shortest and easiest set up time with only 4 XS tasks that were all well documented.

Matillion Data Loader had the longest set up with the most steps. Also, Matillion had some steps that were poorly documented.

Overall, Stitch came out between Fivetran and Matillion Data Loader, but it had the longest-running individual task—selecting which Salesforce entities to sync. All entities were turned “off” by default, and we had to scroll and select each desired one, as well as verify which columns to sync. That task alone took us over 20 minutes to complete.

Test 2: Automation Test Results

After completing the initial setup and sync of data, we made changes to the source data model (added new columns, altered existing columns, etc.) as outlined in the test methodology. To report our results, we constructed another workflow diagram that shows either automation or the effort required to complete manual configuration needed for the data pipeline tool to pick up and start syncing the data source changes.

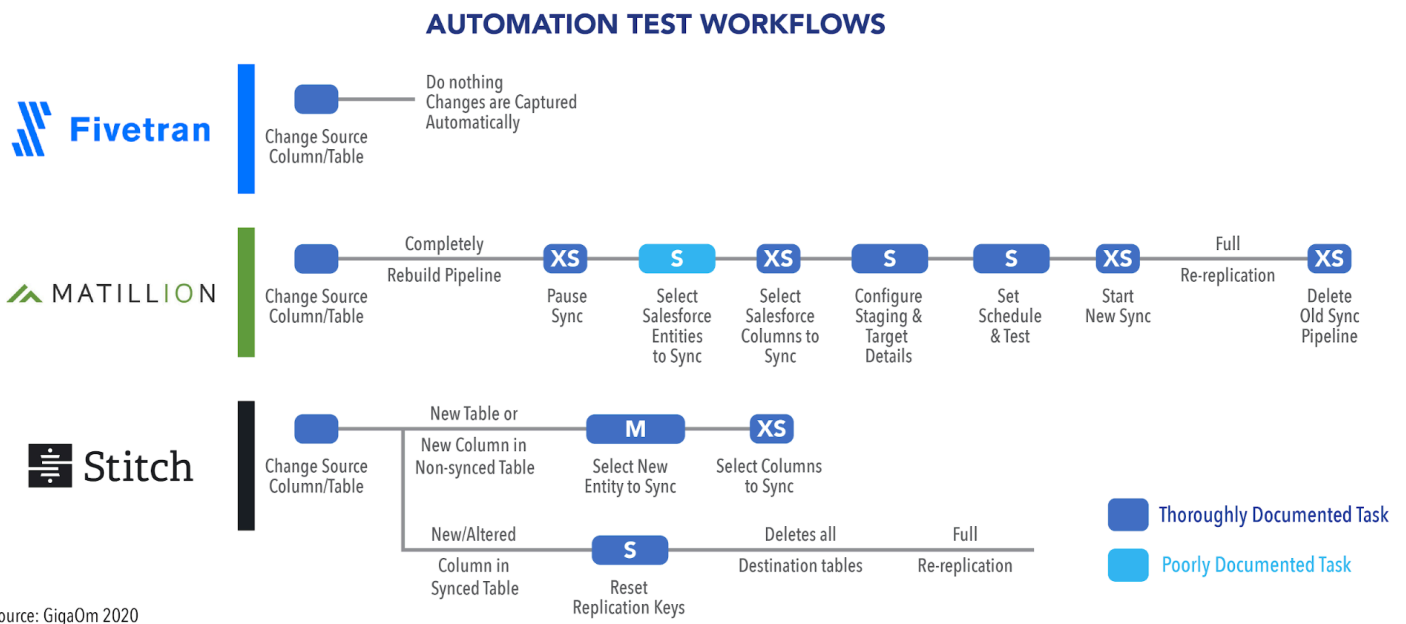


Figure 3. Automation Test Workflows

Workflows represented above have the same characteristics as in the previous diagram, in terms of the width of each task and the level of documentation required to complete tasks.

Fivetran handled the data source changes with full automation. As the user, we had to do nothing to intervene or tell Fivetran to begin sourcing the changes that ranged from updated data to schema changes. The new data/altered columns appeared automatically in our destination data warehouse.

Matillion Data Loader presented the biggest automation challenge, because not only did the new data/ altered columns not appear automatically, but the pipeline had to be rebuilt. This was also confirmed by the Matillion question and answer webpage as a missing feature (dated March 24, 2020):

Q: Is it possible to edit the settings of a pipeline other than the refresh time and notification options? A: We hope to make this possible in a future release. But for now, you can create a new pipeline with new columns and tables, and just delete or disable the old one. ¹

Stitch also required manual intervention to work with new data/altered columns. In the case of a new table or a new column in a previously non-synced table, the pipeline has to be manually altered. The new entity/column added to the pipeline definition. In the event an altered column is in an already synced table, it may be necessary to reset Stitch's Replication Keys. Doing so deletes the data from the destination tables in the warehouse and forces a full re-replication of data. This could have cost impacts, particularly for large tables. Finally, if a new column is added to an already synced table, it was automatically brought over with Stitch.

In terms of automatic configuration and "hands off" experience, only Fivetran has a fully automated data pipeline. Both Matillion and Stitch required some sort of manual intervention or effort whenever a change in the source system occurred.

Test 3: Documentation Test Results

The results of our documentation test are best summarized in the rubric below. However, a few comments should be noted.




Fivetran had the most thorough documentation across all the items we measured, which helps the user understand source set up and contextual source information. Stitch also had good documentation, with only a few items either omitted or left short. Matillion Data Loader's documentation describing the data of the source data connector for Salesforce was nearly completely missing. Also, the reader should note that these patterns of documentation for source connectors extend beyond just their respective Salesforce connectors (at the time of the report, of course).

Scoring Rubric

The following table reveals the complete set of scores we attained during our field test of these data pipeline platforms.

Table 2. Scoring Rubric Results

SCORING RUBRIC RESULTS

	 Fivetran	 MATILLION DATA LOADER	 Stitch
SETUP			
1a. Source Setup Effort	5	3	5
1b. Source Setup Documentation	5	2	5
1c. Destination Setup Effort	5	3	3
1d. Destination Setup Documentation	5	4	5
1. Setup Test Average	5.0	3.0	4.5
INTEGRATION AUTOMATION			
2a. Capture New Column	5	1	3
2b. Capture Renamed Column	5	1	3
2c. Capture Altered Column Data Type	5	1	3
2d. Retain Old Column After Alter/Rename	5	1	5
2e. Capture New Table	5	1	3
2. Integration Automation Test Average	5.0	1.0	3.4
DOCUMENTATION			
3a. Source/Destination Table Relations	5	1	1
3b. Entity Table List	3	1	3
3c. Data Source Usage	5	1	1
3d. Connector Load Behavior	5	3	5
3. Documentation Test Average	4.5	1.8	2.5
OVERALL SCORE	4.8	1.9	3.5

Source: GigaOm 2020

In our field test, Fivetran scored highest with a 4.8. Stitch was second with a 3.5 aggregate score while Matillion Data Loader was third with a 1.9 score.

Again, we understand there is a degree of subjectivity with these scores. We gave our best effort to minimize this issue and focused on specific, measurable attributes to help our readers understand each of these three cloud data pipeline platform's strengths and weaknesses in terms of ease of setup/use, automation, and documentation.


Please note that all three platforms passed the data integrity check. That is, all the data we asked them to sync was fully synced to our destination warehouse. This would be a major problem had a platform failed this check.

Comments on Sync Behavior

While not an actual measurement in our field test, we thought it would be useful to our readers to describe what we learned about each of these three data pipeline tool's syncing behavior. We were interested to know how long it took to perform the initial sync of data and how many queries were run. The reasons have cost implications. Our source system, Salesforce, has a different version or release levels (Professional, Enterprise, Unlimited, etc.) Some of these levels have daily API quotas or a maximum number of API calls that can be made in a day. Each of the three data pipeline tool's Salesforce connectors use the Salesforce API, so the fewer calls it makes, the less likely the data pipeline tool is to exceed your Salesforce daily API quota. To increase your Salesforce API quota costs money, of course.

The following chart shows the number of queries executed by each data pipeline on Snowflake during the initial sync of Salesforce data. Note, however, not every query represents an API call to Salesforce—only a small percentage of them do. Still, it shows the level of activity between Salesforce and Snowflake caused by the data pipeline tool.

NUMBER OF QUERIES EXECUTED DURING THE INITIAL SYNC OF SALESFORCE DATA



Source: GigaOm 2020

Figure 4. Number of Queries Executed During the Initial Sync of Salesforce Data

Additionally, each platform executed a variety of queries against Snowflake during the initial sync of Salesforce data. The following diagram shows the volume and variety of the different query types. As you can see, Fivetran executed the fewest different types of queries. Stitch executed the widest variety. Matillion executed an average variety of queries--it just executed a lot of them.

VARIETY OF QUERIES EXECUTED DURING THE INITIAL SYNC OF SALESFORCE DATA


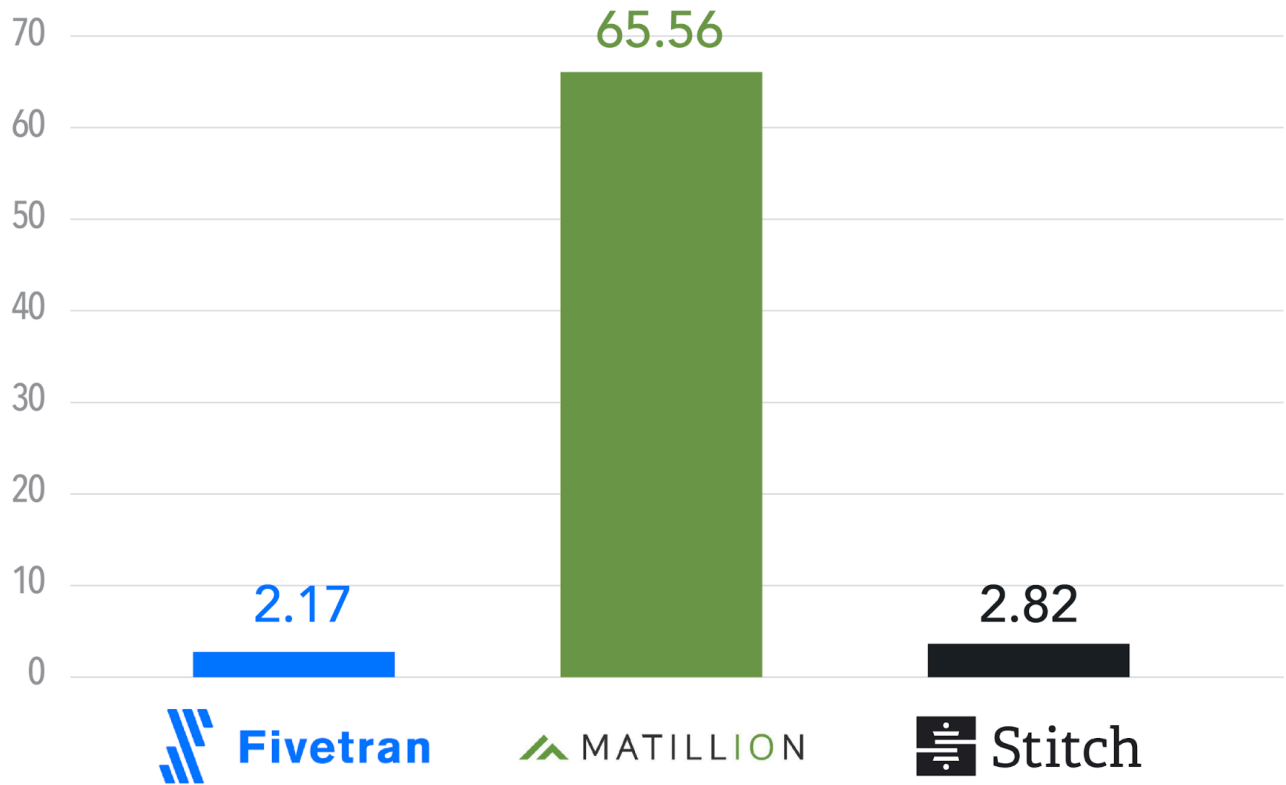


Figure 5. Variety of Queries Executed During the Initial Sync of Salesforce Data

Our destination data warehouse, Snowflake, charges for the elapsed time it runs. Thus, if the data pipeline tools run more queries or take longer to sync the same amount of data, that will potentially increase your Snowflake costs.

The following chart shows the number of Snowflake credits used by each platform to sync the initial Salesforce data set of 10,000 Accounts, 50,000 Contacts, and 50,000 Leads, plus all the ancillary data that comes with Salesforce.

During our test, Fivetran ran the fewest number of queries and used the fewest number of Snowflake credits compared to Matillion or Stitch.

SNOWFLAKE CREDITS USED TO SYNC INITIAL SALESFORCE DATA

Source: GigaOm 2020

Figure 6. Snowflake Credits Used to Sync Initial Salesforce Data

Snowflake credits are then multiplied by a dollar amount. This often amounts to \$2.00, \$3.00, or \$4.00 per credit (depending on the type of Snowflake account you have). While these were not expensive syncs, you can see how the sync activity will add up over time.

Again, these are not “official” results of our Field Test, and will vary widely depending on your data and use case. However, as a baseline observation, we found these observations compelling enough to share.

¹ <https://www.matillion.com/resources/blog/matillion-data-loader-12-questions-answers/>

5. Conclusion

In our test, Fivetran had the shortest and easiest setup. Matillion Data Loader had the longest setup with the most steps. Also, Matillion had some steps that were poorly documented. Stitch was in between Fivetran and Matillion Data Loader, but it had the longest-running individual task when selecting which Salesforce entities to sync.

Fivetran handled the data source changes with full automation. Matillion Data Loader presented the biggest automation challenge, because not only did the new data/altered columns not appear automatically, but the pipeline had to be rebuilt. Stitch also had manual interventions that need to take place in the event of new data/altered columns.

Fivetran had the most thorough documentation across all the items we measured. Stitch also had good documentation, with only a few items either omitted or left short. Matillion Data Loader's documentation describing the data of the source data connector for Salesforce was nearly completely missing.

The level of activity between Salesforce and Snowflake caused by Matillion was excessive compared to both Stitch and Fivetran.

These tests are important because while our testing shows just one source to one destination, there will be a multiplier effect depending on how many connections your company requires.

Automated data integration is well worth exploring for any enterprise data integration need, especially where your source and target are supported.

6. Disclaimer

This test is a point-in-time check into specific attributes of these tools. There are numerous other factors to consider in the selection process, ranging across Administration, Integration, Workload Management, User Interface, Scalability, Vendor, Reliability, and numerous other criteria. It is also our experience that product characteristics change over time and are competitively different for different workloads. Also, a leader can run up against the point of diminishing returns and viable contenders can quickly close the gap.

GigaOm runs all of its comparison tests according to strict ethical standards. The results of the report are the objective results of the application of queries to the simulations described in the report. The report clearly defines the selected criteria and process used to establish the field test. The report also clearly states the data set sizes, the platforms, the configurations, etc. used. The reader is left to determine for themselves how to qualify the information for their individual needs. The report does not make any claim regarding the third-party certification and presents the objective results received from the application of the process to the criteria as described in the report. The report strictly measures the attributes indicated and does not purport to evaluate other factors that potential customers may find relevant when making a purchase decision.

This is a sponsored report. Fivetran chose the competitors, the test, and the Fivetran configuration. GigaOm chose the most compatible configurations for the other tested platform and ran the testing workloads. Choosing compatible configurations is subject to judgment calls. We have attempted to describe our decisions in this paper.

7. About Fivetran

Fivetran, the leader in automated data integration, delivers ready-to-use connectors that automatically adapt as schemas and APIs change, ensuring consistent, reliable access to data. Fivetran improves the accuracy of data-driven decisions by continuously synchronizing data from source applications to any destination, allowing analysts to work with the freshest possible data. To accelerate analytics, Fivetran enables in-warehouse transformations and delivers source-specific analytics templates. With more than 1,000 customers, Fivetran is headquartered in Oakland, California, with offices around the globe. For more information, visit www.fivetran.com.

8. About William McKnight



An Ernst & Young Entrepreneur of the Year Finalist and frequent best practices judge, William is a former Fortune 50 technology executive and database engineer. He provides Enterprise clients with action plans, architectures, strategies, and technology tool selection to manage information.

William McKnight is an Analyst for GigaOm Research who takes corporate information and turns it into a bottom-line producing asset. He's worked with companies like Dong Energy, France Telecom, Pfizer, Samba Bank, ScotiaBank, Teva Pharmaceuticals and Verizon — Many of the Global

2000 — and many others. William focuses on delivering business value and solving business problems utilizing proven, streamlined approaches in information management.

He is a frequent international keynote speaker and trainer. William has taught at Santa Clara University, UC-Berkeley and UC-Santa Cruz.

9. About Jake Dolezal



As a contributing Analyst at GigaOm, Jake Dolezal has two decades of experience in the Information Management field with expertise in analytics, data warehousing, master data management, data governance, business intelligence, statistics, data modeling and integration, and visualization. Jake has experience across a broad array of industries, including: healthcare, education, government, manufacturing, engineering, hospitality, and restaurants. He has a doctorate in information management from Syracuse University.

10. About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

11. Copyright

© [Knowingly, Inc.](#) 2020 "*Data Pipeline Platform Comparison*" is a trademark of [Knowingly, Inc.](#). For permission to reproduce this report, please contact sales@gigaom.com.