



**SENTURUS**

# **Azure Synapse vs. Snowflake Comparing Analytic Data Clouds**

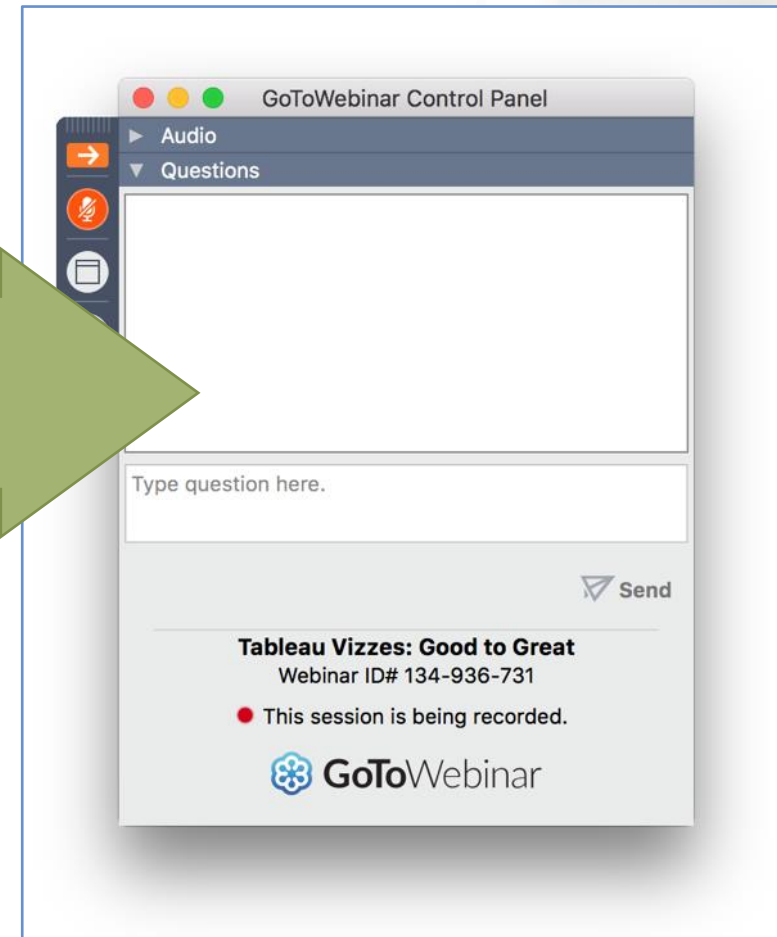


# GoToWebinar control panel

Click arrow to restore full control panel



Submit questions in this section



# To obtain this presentation

Visit the Resource Library  
on the Senturus website  
to download this presentation  
and explore other assets

[senturus.com/resources](https://senturus.com/resources)

## RESOURCE LIBRARY

Welcome to our extensive, free library of past webinars, demos, whitepapers, presentations and helpful hints. Use the topic boxes to the right to filter through and easily locate content. We are constantly adding new materials, so please check back often to see our latest content. Enjoy!

SEARCH RESOURCES

GO

SORT BY

Popularity

Date

Title

Type



RESOURCE TITLE

TYPE

GO

SHEDDING LIGHT ON CLOUD BI OPTIONS  
Amazon. IBM Cognos Cloud. Microsoft Power BI Service. Tableau Online.  
Azure.



BUSINESS  
STRATEGY



IBM COGNOS REPORTING TOOLS: FEATURE SET COMPARISON  
Visual Reference Guide



FEATURES



A COMPARISON OF POWER BI, TABLEAU & COGNOS  
Differentiators Demo'd



FEATURES



# Agenda

Introduction

Azure Synapse vs Snowflake Comparisons

Senturus overview

Additional resources

Q&A





# Introductions



**Bob Looney**

VP Software & Architecture  
Senturus, Inc.



**Reeves Smith**

Principal Snowflake Architect  
Senturus, Inc.



**Gary Lock**

Sr. Microsoft Architect  
Senturus, Inc.



**Michael Weinbauer**

Director  
Senturus, Inc.

# Poll

Where is your data stored now?

- On-prem
- VM based (lift and shift) cloud
- Cloud native app (Snowflake, Azure Synapse)

# Azure Synapse vs. Snowflake Comparisons

- Platform architecture
- Tooling
- PaaS vs. SaaS
- Compute architecture
- Azure specific features
- Snowflake specific features
- Pricing

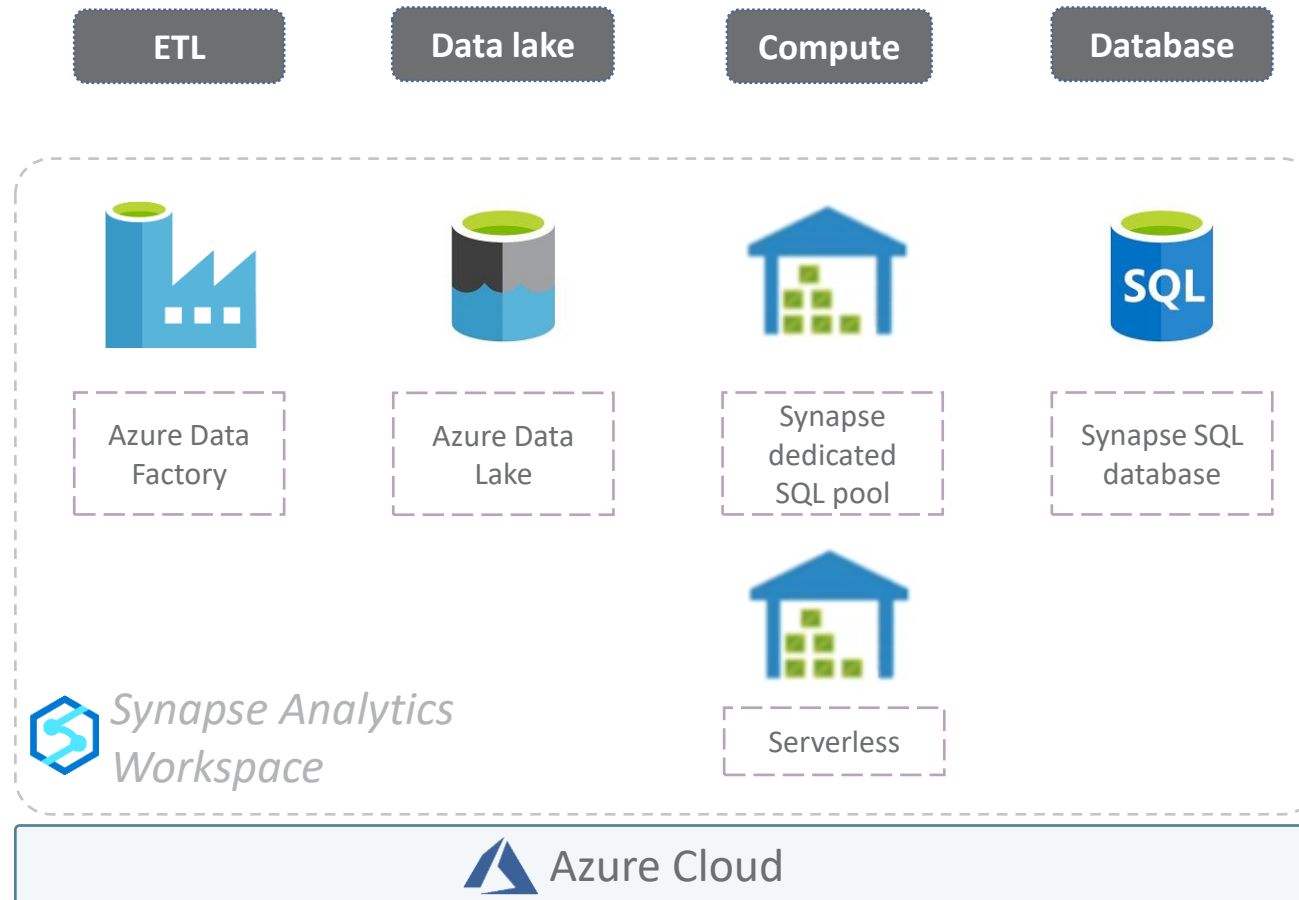
# Platform architecture

- Cloud-based
- Highly scalable storage for data files
- Demand-based computing resources
- SQL database for curated data storage
  - Columnar approach to storing data, ideal for analytics
  - Designed with massive concurrency in mind
- Covers modern audiences and use cases
  - Classic BI
  - Self service
  - Data science / AI

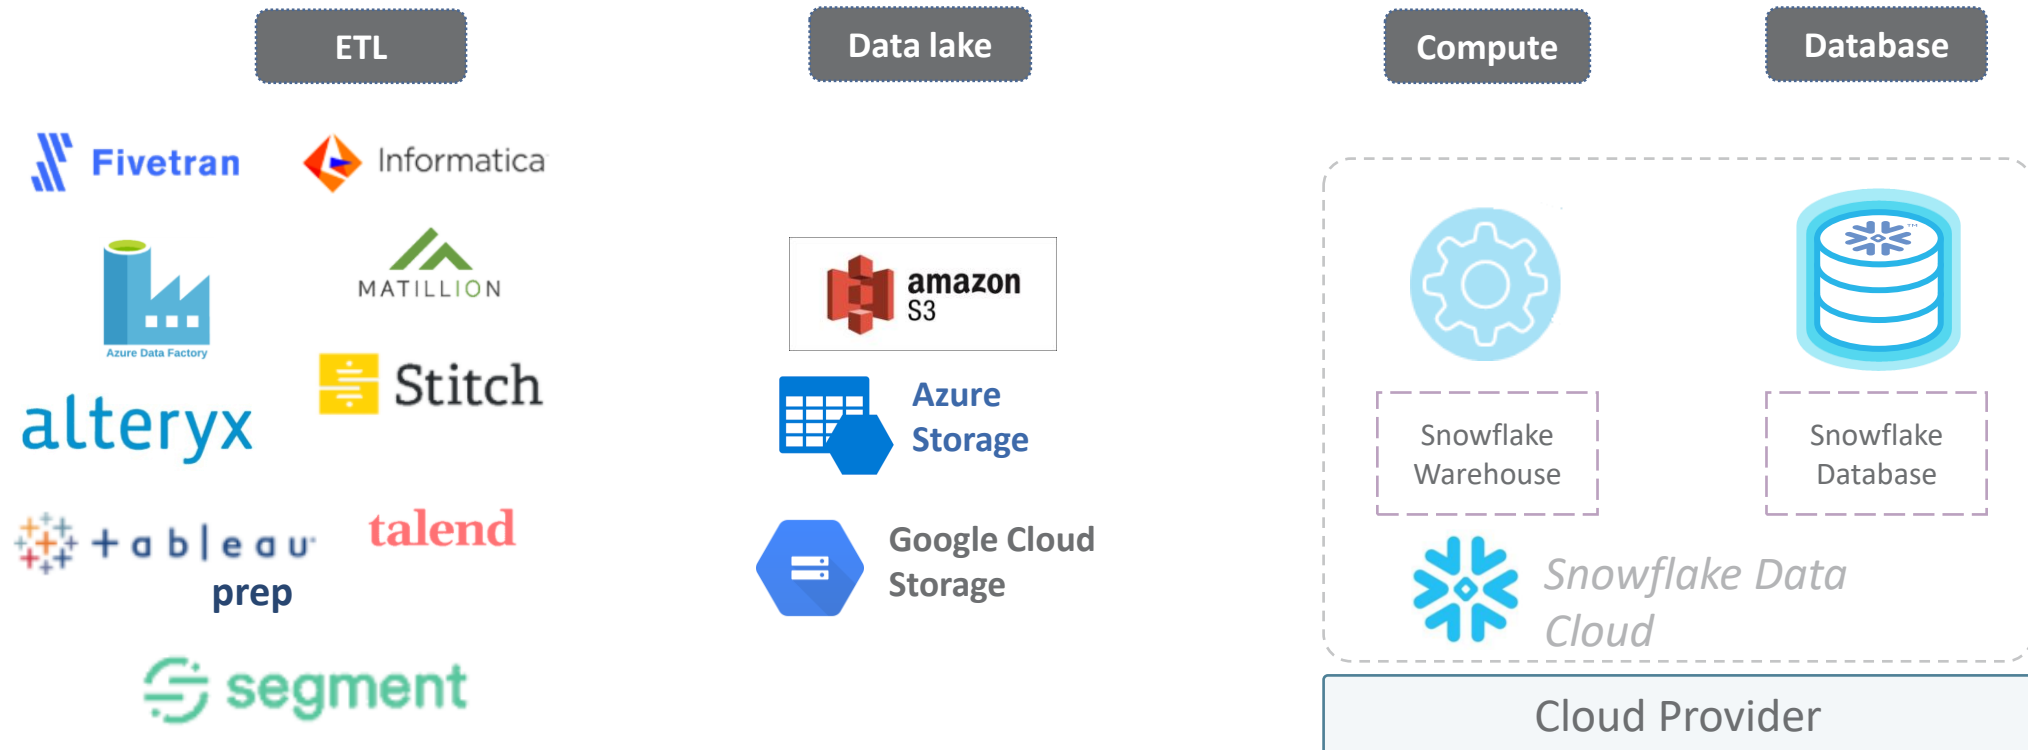




# Azure Synapse architecture



# Snowflake architecture



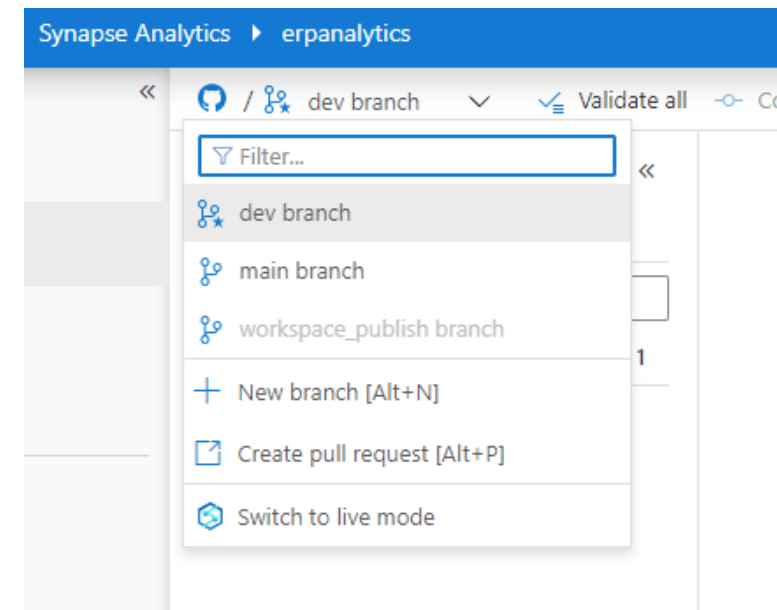
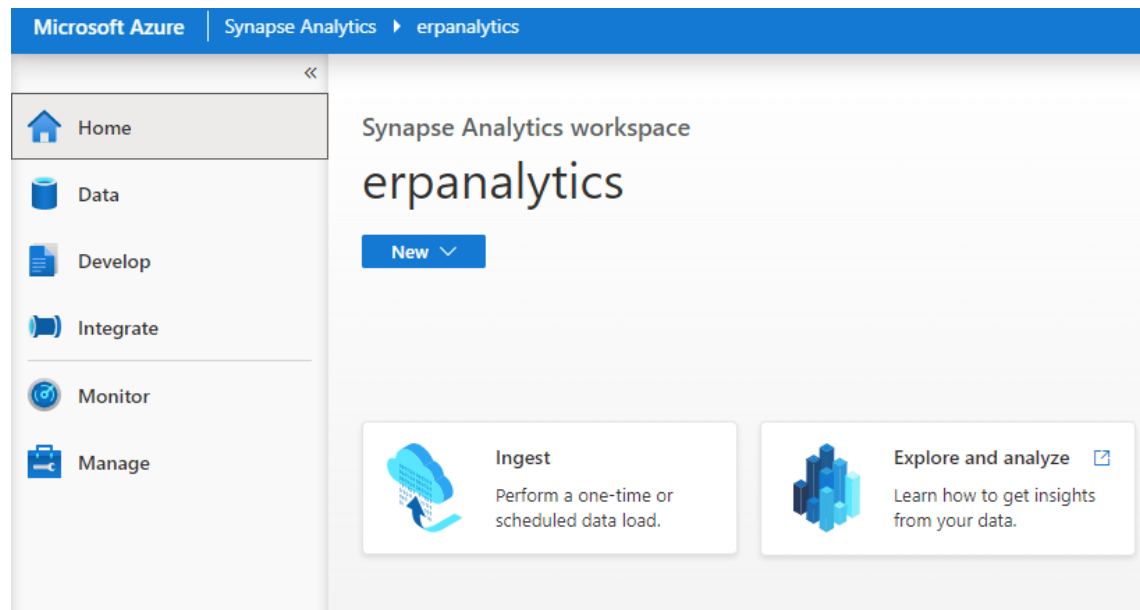
# Wide, robust tooling

- Both provide web interfaces for management and development tasks
- Support from major 3<sup>rd</sup> party data analytics and ETL tools
- ODBC drivers, command line tools



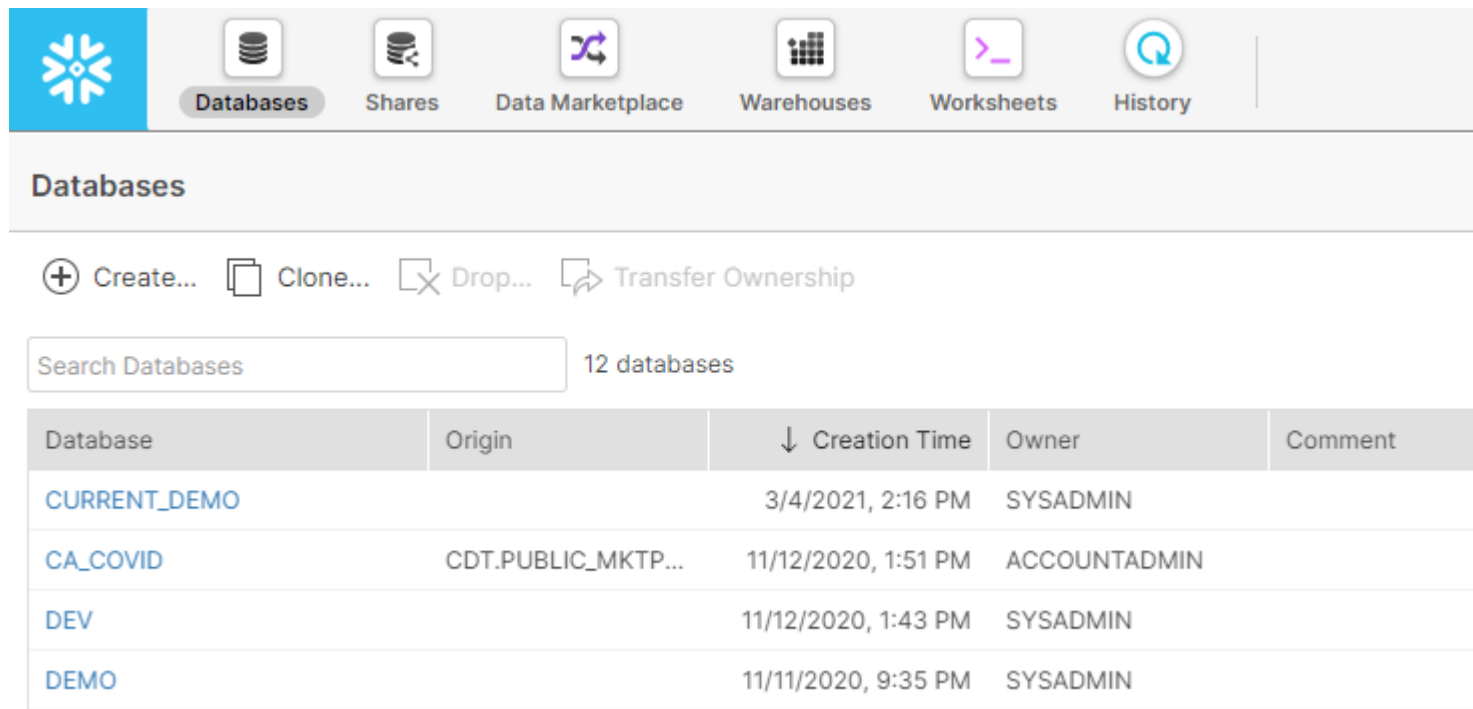
# Azure Synapse

- Synapse workspace web development interface has seen a lot of recent investment and is becoming full featured
- SSMS is compatible with Synapse SQL databases



# Snowflake

- Web interface robust for compute and database components
- Tools like DataGrip can be used to create a more robust development environment

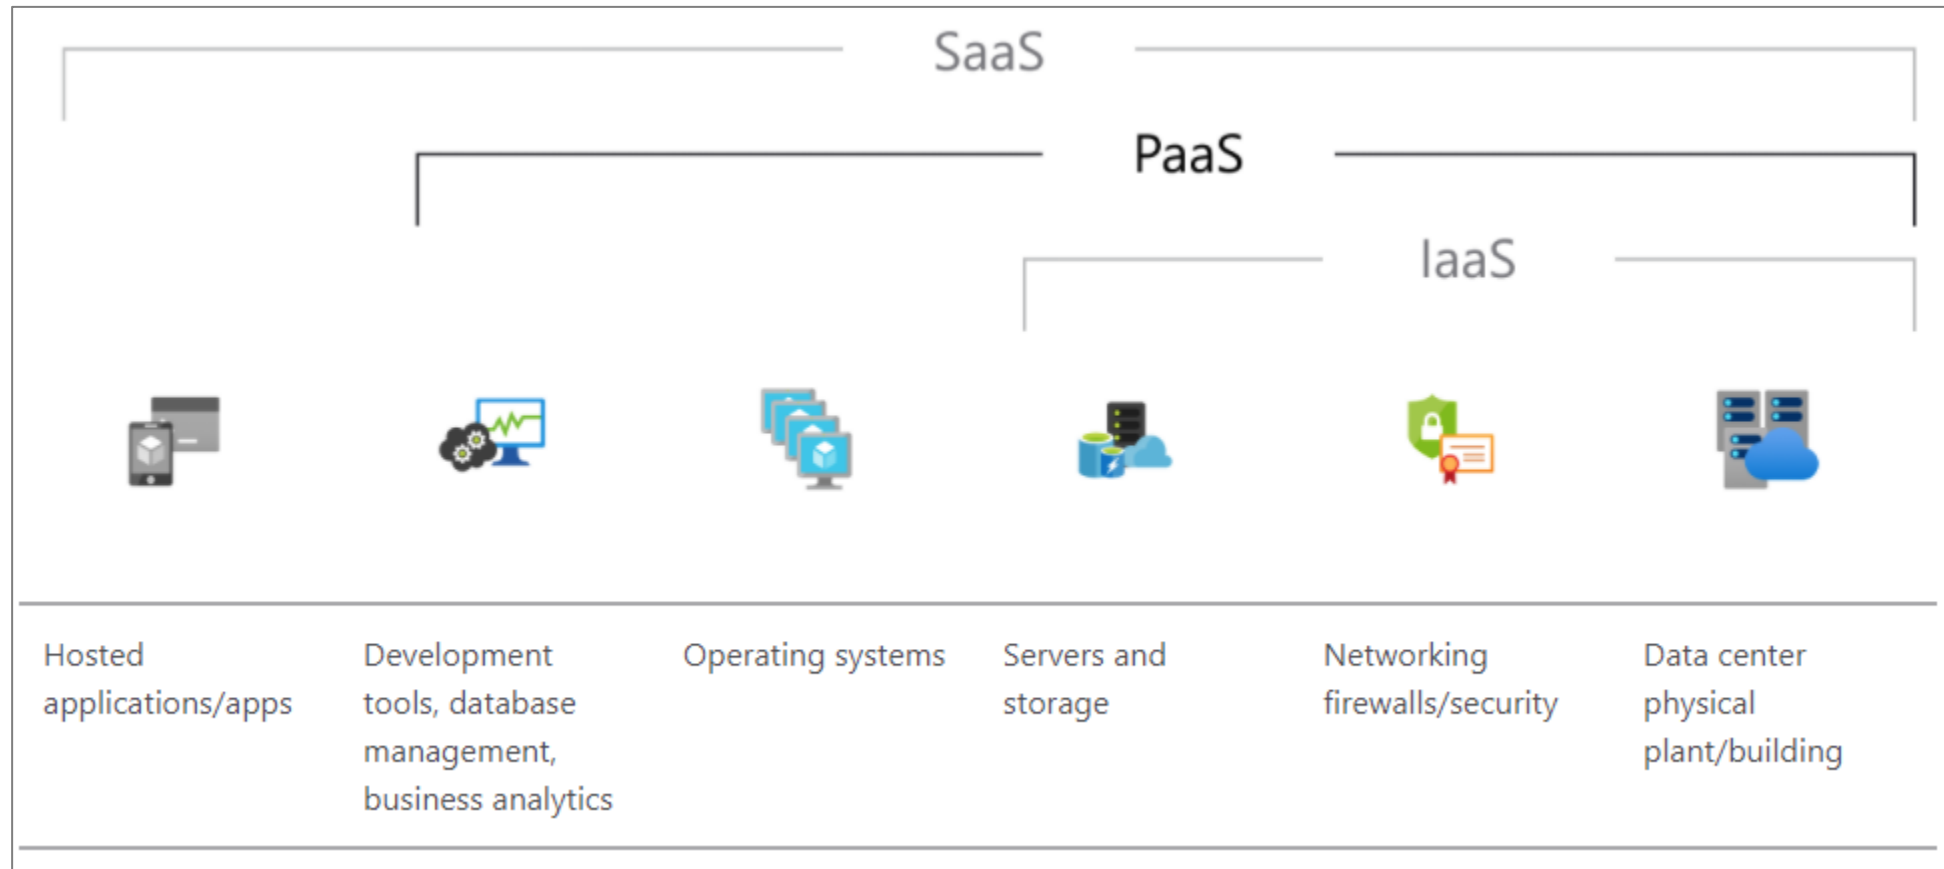


The screenshot displays the Snowflake web interface. At the top, there is a navigation bar with icons for Databases, Shares, Data Marketplace, Warehouses, Worksheets, and History. The 'Databases' tab is selected. Below the navigation bar, the 'Databases' section is visible. It includes a search bar labeled 'Search Databases' and a count of '12 databases'. Below the search bar, there is a table listing databases. The table has columns for Database, Origin, Creation Time, Owner, and Comment. The databases listed are CURRENT\_DEMO, CA\_COVID, DEV, and DEMO.

Database	Origin	Creation Time	Owner	Comment
CURRENT_DEMO		3/4/2021, 2:16 PM	SYSADMIN	
CA_COVID	CDT.PUBLIC_MKTP...	11/12/2020, 1:51 PM	ACCOUNTADMIN	
DEV		11/12/2020, 1:43 PM	SYSADMIN	
DEMO		11/11/2020, 9:35 PM	SYSADMIN	



# PaaS vs. SaaS



# PaaS vs. SaaS – Azure Synapse

## Synapse - platform as a service

- A collection of Azure resources
- Data Lake Gen2, Azure SQL Server (custom), Azure Data Factory
- Created and maintained in Azure portal
- Synapse workspace web interface consolidated view

Microsoft Azure

Search resources, services, and docs (G+)

Home > Azure Synapse Analytics >

**Azure Synapse Ana...** «

Senturus, Inc.

+ Create ⚙ Manage view ▾ ...

Filter for any field...

Name ↑↓

erpanalytics	...
--------------	-----

**Create Synapse workspace** ...

\* Basics \* Security Networking Tags Review + create

Create a Synapse workspace to develop an enterprise analytics solution in just a few clicks.

**Project details**

Select the subscription to manage deployed resources and costs. Use resource groups like f of your resources.

Subscription \* ⓘ CSP Subscription

Resource group \* ⓘ

Managed resource group ⓘ Enter managed resource group name

Create new

# PaaS vs. SaaS - Snowflake

## Snowflake - software as a service

- Snowflake has an abstraction layer between the actual cloud resources
- This allows the same Snowflake software to run on multiple clouds - AWS, Azure and GCP clouds
- Created at snowflake.com, login to tenant on selected cloud and region



# Compute architecture

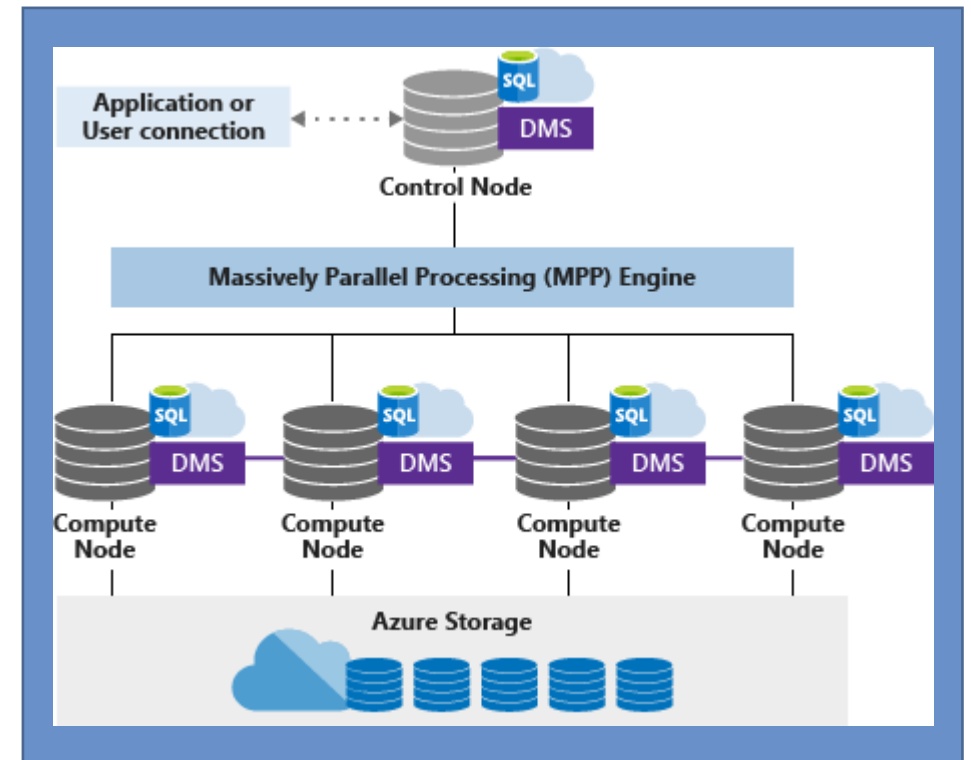
When performing actions like data loading or querying, compute capacity is needed.

Capacity is priced by how long it is run and what performance level is selected.



# Compute architecture – Azure Synapse

- Dedicated SQL pool aligns with 1 SQL database
- Database can only be accessed when dedicated SQL pool is running
- Massively parallel under the covers
- Workload groups and classifiers to isolate and prioritize actions
- Manual or API scale, stop, resume





# Compute architecture - Snowflake

- Compute (“warehouse”) is fully separate from any Snowflake database
- Multiple compute can access the same database simultaneously
- The same compute can access multiple databases
- Manual scale up/down
- Auto scale in/out, stop, resume



# Azure Synapse specific

## Apache Spark pools

- Prep data and build machine learning models (python, scala, .NET) using notebooks in the web development environment

The screenshot displays the Azure Synapse Studio interface. On the left, the 'Data' pane shows a workspace with a search bar 'Filter resources by name'. Under 'Databases', there is one database named 'webinardw (SQL)'. Under 'Tables', there is one table named 'dbo.Webinar'. Below this, there are sections for 'External tables' and 'External resources'. The main area is titled 'Notebook 1' and contains a toolbar with options like 'Cell', 'Run all', 'Publish', 'Attach to' (set to 'Select Apache Spark pool'), and 'Language' (set to 'Spark (Scala)'). A message states 'Not started' and 'Please select a Spark pool to attach before running cell!'. Below this, a code cell is shown with the following Scala code:

```
1 %%spark
2 val df = spark.read.sqlanalytics("webinardw.dbo.Webinar")
3 df.write.mode("overwrite").saveAsTable("default.t1")
```


A context menu is open over the code cell, showing options: 'New SQL script', 'New notebook', and 'Load to DataFrame'.

# Snowflake specific

## Zero-copy data cloning

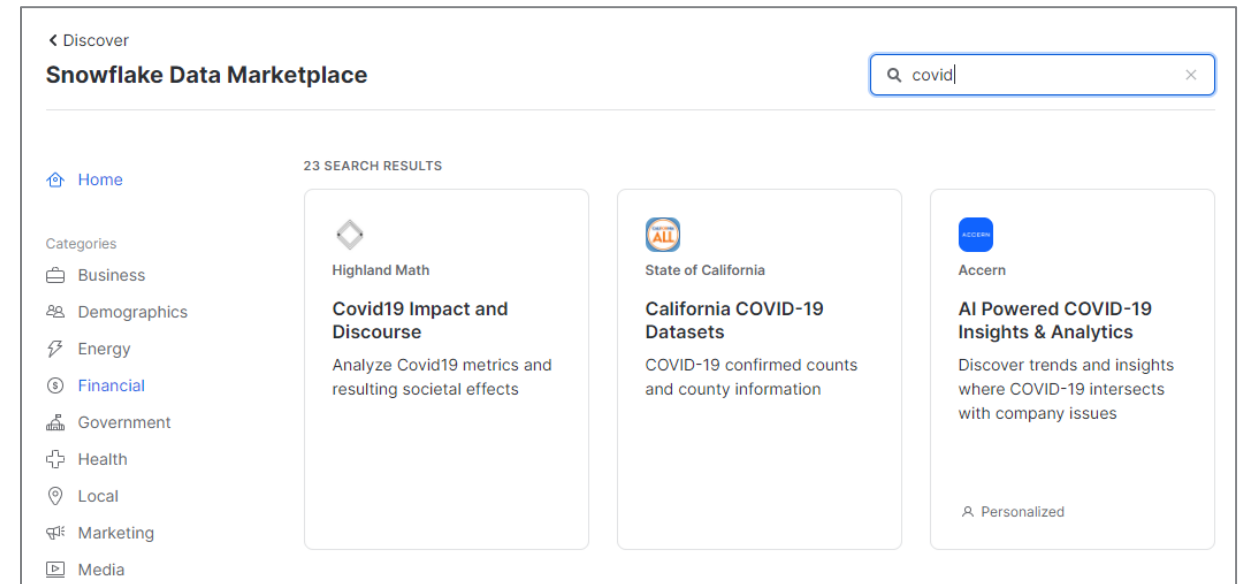
- Major benefits for DevOps

```
CREATE DATABASE Dev CLONE Prod;
```



## Data sharing

- Pull in data from the Marketplace
- Share data securely and easily with partners



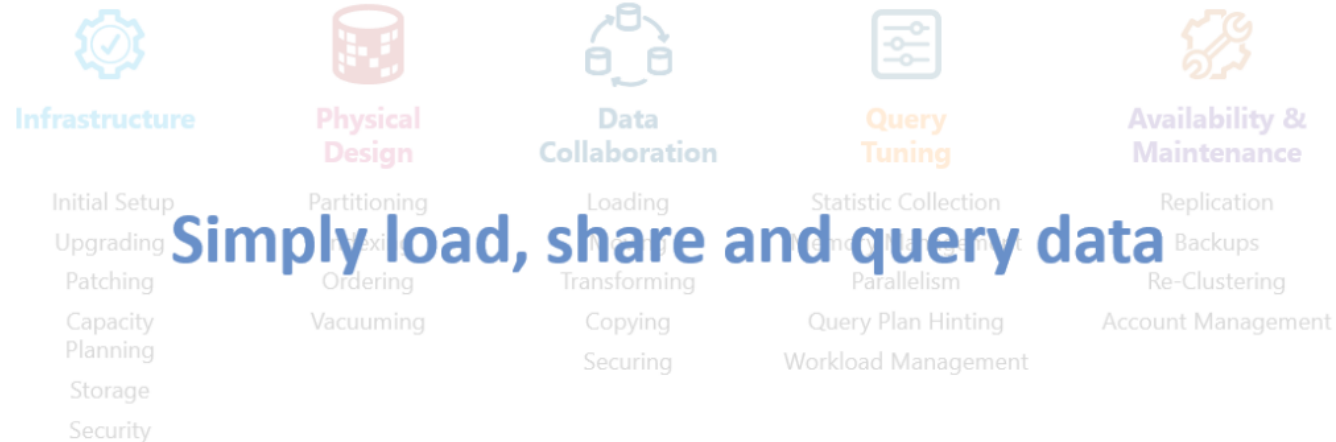
# Snowflake specific

## Time Travel queries

```
select * from my_table at(offset => -60*5);
```

## Less administration

### SNOWFLAKE CLOUD DATA PLATFORM



Webinar: 10 Reasons  
Snowflake is Great  
for Analytics

# Pricing

- Compute and storage are billed separately
- Consumption-based
- Compute can be scaled, paused and resumed
- Reserved or pre-purchased can yield savings





# Azure Synapse pricing

- Compute pricing based on the power of the dedicated SQL pool
- SQL database storage
- Raw file storage pricing based on Azure data lake Gen2
- Serverless \$5/TB processed

SERVICE LEVEL	DWU	PRICE
DW100c	100	\$1.20/hour
DW200c	200	\$2.40/hour
DW300c	300	\$3.60/hour
DW400c	400	\$4.80/hour
DW500c	500	\$6/hour
DW1000c	1000	\$12/hour

Data storage is charged at the rate of \$23 per TB of data processed (\$0.04/1 TB/hour). Data storage includes the size of your data warehouse and 7-days of incremental snapshot storage.

Data storage prices		
	PREMIUM	HOT
First 50 terabyte (TB) / month	\$0.15 per GB	\$0.0184 per GB
Next 450 TB / month	\$0.15 per GB	\$0.0177 per GB
Over 500 TB / month	\$0.15 per GB	\$0.0169 per GB

# Snowflake pricing

- Compute pricing based on the size of the warehouse and edition
- SQL database storage
- Data lake raw file storage is separate

\$40 per Terabyte per month for customers deployed in Azure – West US 2 (Washington)	
ON DEMAND	Price per Snowflake Credit
Standard	\$2.00
Enterprise	\$3.00
Business Critical	\$4.00

# Constant use pricing comparison

- Assume a 1TB database and lowest compute capacity running 24x7 for 30 days and “on demand” retail pricing
- Ignoring ETL, file storage and data transfer (assume equivalent)

## Azure Synapse dedicated

- DW100c = \$1.20/hr
- 1TB SQL = ~\$30/month

**Total = \$894/month**

## Snowflake

- xs warehouse = \$2.00/hr
- 1TB SQL = \$40/month

**Total = \$1,480/month**

# Variable demand pricing comparison

- Assume a 1TB database and lowest compute capacity running a business hours for 4 weeks and “on demand” retail pricing
- Ignoring ETL, file storage and data transfer (assume equivalent)

## Azure Synapse dedicated

- DW100c = \$1.20/hr
  - 10 hours, 5 days/week = 200 hrs
- 1TB SQL = ~\$30/month

**Total = \$270/month**

## Snowflake

- xs warehouse = \$2.00/hr
- Auto-pause and start = 50 hrs
- 1TB SQL = \$40/month

**Total = \$140/month**

# Comparison strengths to consider

	Azure Synapse	Snowflake
Ease of use		X
Vendor provided tooling	X	
Ease of administration		X
Pricing advantage	Constant use	Variable use
O365 & Azure shops	X	
AWS, GCP & other clouds		X
Multiple databases		X
SQL stored procedures	X	*



# Which is the best data platform for *you*?

Turn to Senturus to

Evaluate your unique environment and requirements

Provide an unbiased recommendation (cloud, hybrid, on-prem)

Provide best practices to optimize spend and performance

Guide and/or execute the implementation

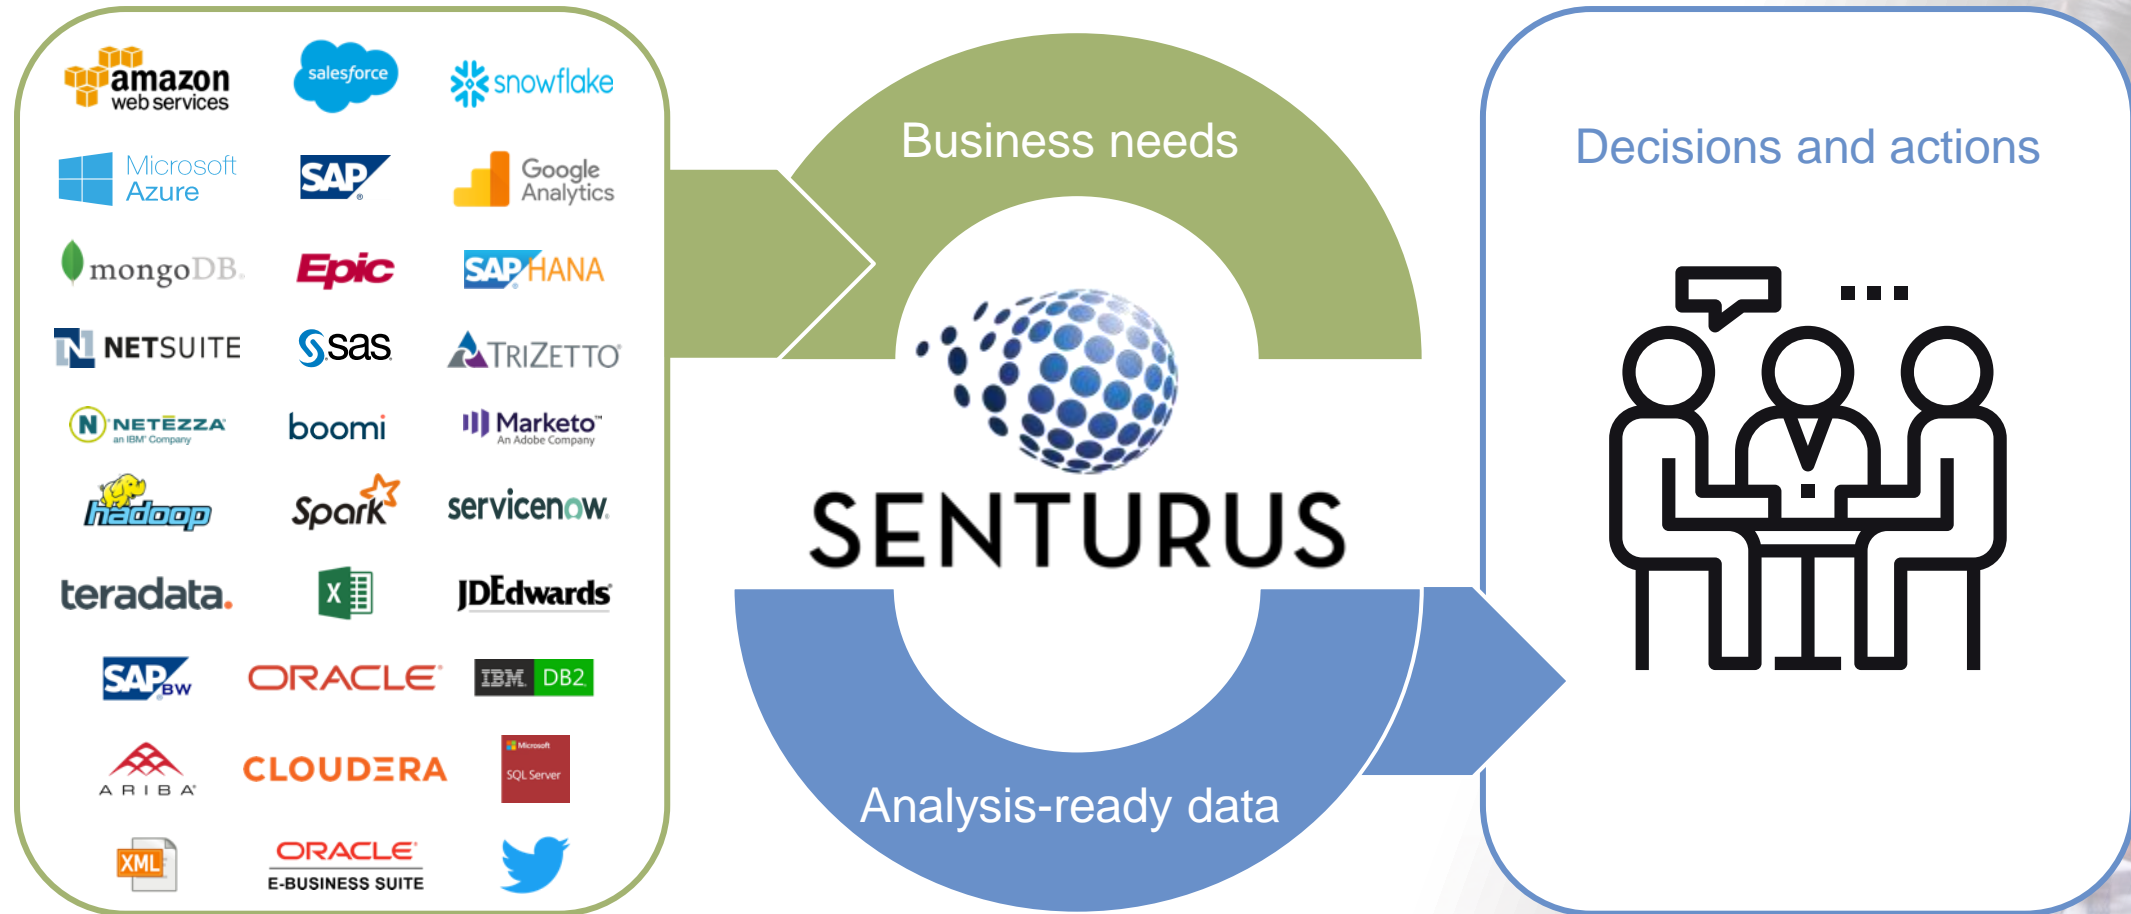
Contact us at [info@Senturus.com](mailto:info@Senturus.com) | 888 601 6010

# The authority in Business Intelligence

Exclusively focused on BI,  
Senturus is unrivaled in its  
expertise across the BI stack



# Bridging the gap



# Full spectrum of BI services

- Dashboards, reporting and visualizations
- Data preparation and modern data warehousing
- Hybrid BI environments (migrations, security, etc.)
- Software to enable bimodal BI and platform migrations
- BI support retainer (expertise on demand)
- Training and mentoring



# A long, strong history of success

- 19+ years
- 1300+ clients
- 2500+ projects



# Expand your knowledge

Find more resources  
on the Senturus website

[senturus.com/senturus-resources](https://senturus.com/senturus-resources)



# Upcoming events

## **Power Query Editor Jumpstart**

Organizing data to work for your reporting needs

Wednesday, Mar. 17, 2021, 11am PT/2pm ET

## **6 Ways to Publish & Share Tableau Dashboards**

How-to demo and discussion of pros, cons and considerations

Thursday, Apr. 8, 2021, 11am PT/2pm ET



# Complete BI training offerings



**Power BI**



Tailored group sessions




Mentoring



Instructor-led online courses



Self-paced learning



# Additional resources from Senturus

## Unbiased product reviews



## Technical tips



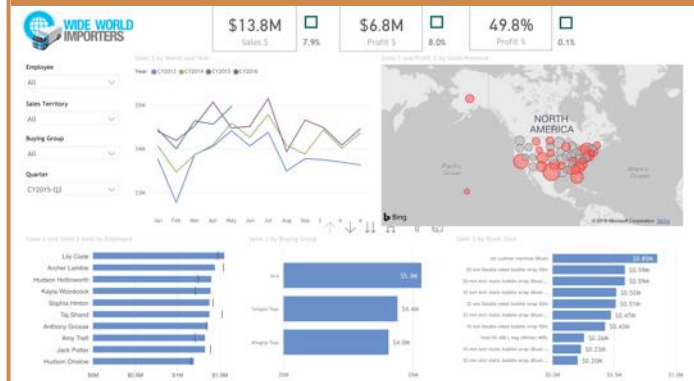
## Insider viewpoints



## More on this subject



## Product demos



## Upcoming events




# Q & A

If your question or issue is broader than what we are able to answer today, contact us at

[info@senturus.com](mailto:info@senturus.com)

and we will set up a free consultation.



# Thank You



## SENTURUS

---

[www.senturus.com](http://www.senturus.com)

888 601 6010

[info@senturus.com](mailto:info@senturus.com)

© 2021 by Senturus, Inc. This presentation may not be reused or distributed without the written consent of Senturus, Inc.