



WAIT! WAIT! BEFORE YOU SIGN UP FOR THAT FREE TRIAL

Laura Ramsey

Director, Technology Evangelism

Snowflake Inc.

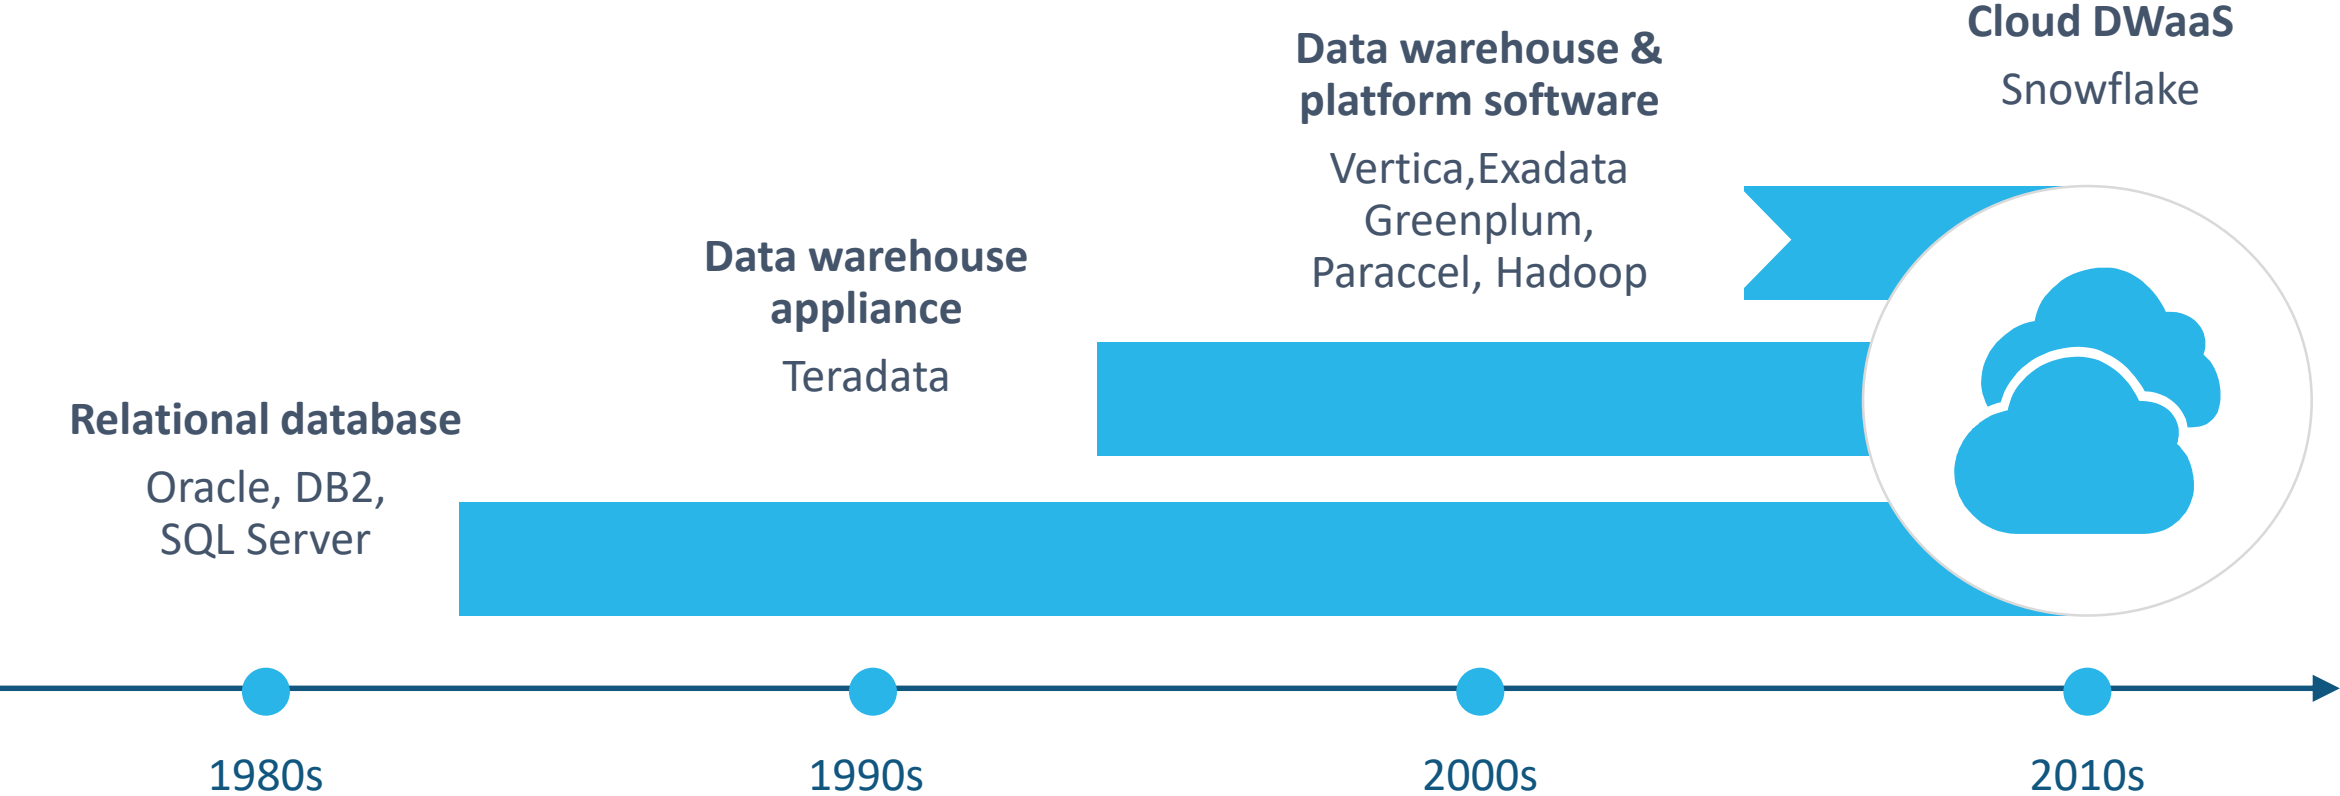
My Bio



- Director, Technical Evangelism @ Snowflake Inc.
- OTN Database Community Manager
- Member of OpenSolaris and OpenJDK Communities (Alumni)
- Technology enthusiast since fax machines...
 - 30+ years in Fortune 50 Tech Companies
- Responsible for the FIRST viral marketing campaign...



The evolution of data platforms



CLOUD DATA WAREHOUSE POC

Agenda

Defining a Valuable POC

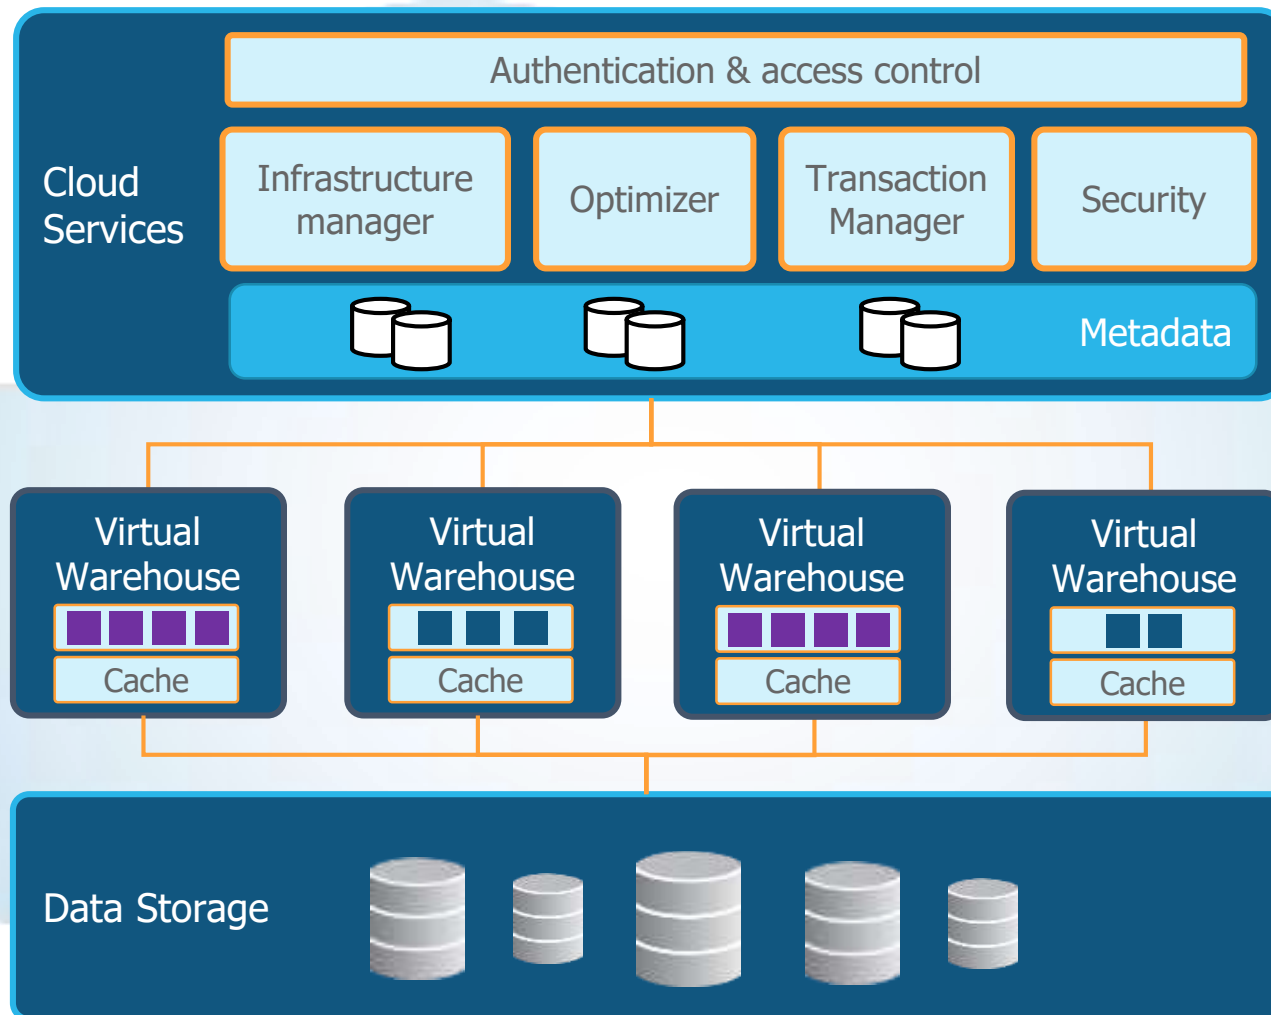
Identifying Capabilities vs. Hype

Setting Expectations

Preparing Data and Applications for Trial

Evaluation Report Card

What are you Test Driving?



Dynamically Combines Three Principal Architectural Layers

- Collection of Independent, Scalable Services
 - “Brain” of the System
 - Handles Crucial Data Management Functionalities
- Elastic Virtual Warehouse
 - “Muscle” of the system
 - Handles Parallel Query Processing
- Storage Layer
 - Hybrid Columnar

WHAT CONSTITUTES A VALUABLE POC?

Prioritizing your technology evaluation scenario

- > What do you need to test drive?
 - > Reporting / analytics
- > What do you want to improve?
 - > Appliance-based data warehouse
- > What are you willing to live with?
 - > Data marts / Silos
- > What can't you live without?
 - > Stored procedures, connectors, access, integration + compatibilities



EVALUATION SCOPE AND REQUIREMENTS

Know it when you see it: Establish a Future State

- Have business goals for the evaluation
- Create vision statement for successful evaluation
- Look for the value (capabilities/change=value)
- Scale evaluation to valuable implementation window



PREPARING FOR THE EVALUATION

Discover, Scope and Document

- Prepare a sample set of databases to migrate
- List database objects to migrate
- Understand the processes and tools that will populate and pull data from other data warehouses
- Have sample roles, users and permissions
- List CloudDW accounts that exist or need creating
- Documentation of the existing data warehouse solution into an “as-is” architecture diagram



PREPARING FOR THE EVALUATION

The Approach

- List as-is
- List re-engineering
- List fixing
- List preferred outcomes/thresholds



PREPARING FOR THE EVALUATION

Define the Evaluation Timeline

- List business expectations for the evaluation
- Document budget allocated for the project*
- Time estimates and staffing / resources

*FREE AS IN BEER....



PREPARING FOR THE EVALUATION

Organize Sample Data Sets

- List data sets to upload first
- Create method for identifying process dependencies for data sets
- Document process dependencies for data sets




PREPARING FOR THE EVALUATION

Determine the Evaluation Outcomes

- List high-level outcomes once the migration is completed
- Document plan for communicating the migration project wins to stakeholders



SECURITY SERVICE CAPABILITIES



Access

- All communication secured & encrypted
- TLS 1.2 encryption in both trusted and untrusted networks
- IP whitelisting
- PrivateLink




Authentication

- Password Policy enforcement
- Multi-factor authentication
- Federated authentication (SAML 2.0)



Authorization/Access Control

- Flexible user management
- Role-based access control for granular control
- RBAC for data and actions



Data

- Encrypted at rest
- Hierarchical key model
- Keys stored in HSM
- Time Travel 1-90 days
- Periodic account-level re-keying of data
- Tri-Secret Secure
- Query statement encryption



Infrastructure

- Physical Security
- Redundancy



EXECUTING THE EVALUATION

Evaluate Security Capabilities Using Scenarios

- Create roles for the first data sets and assign users to roles based on the work they do for the migration
- Establish roles for developer access on non-production databases and read only access, read and write access, and administrative access on all databases
- Develop an automated process for managing Snowflake users and roles early in the evaluation

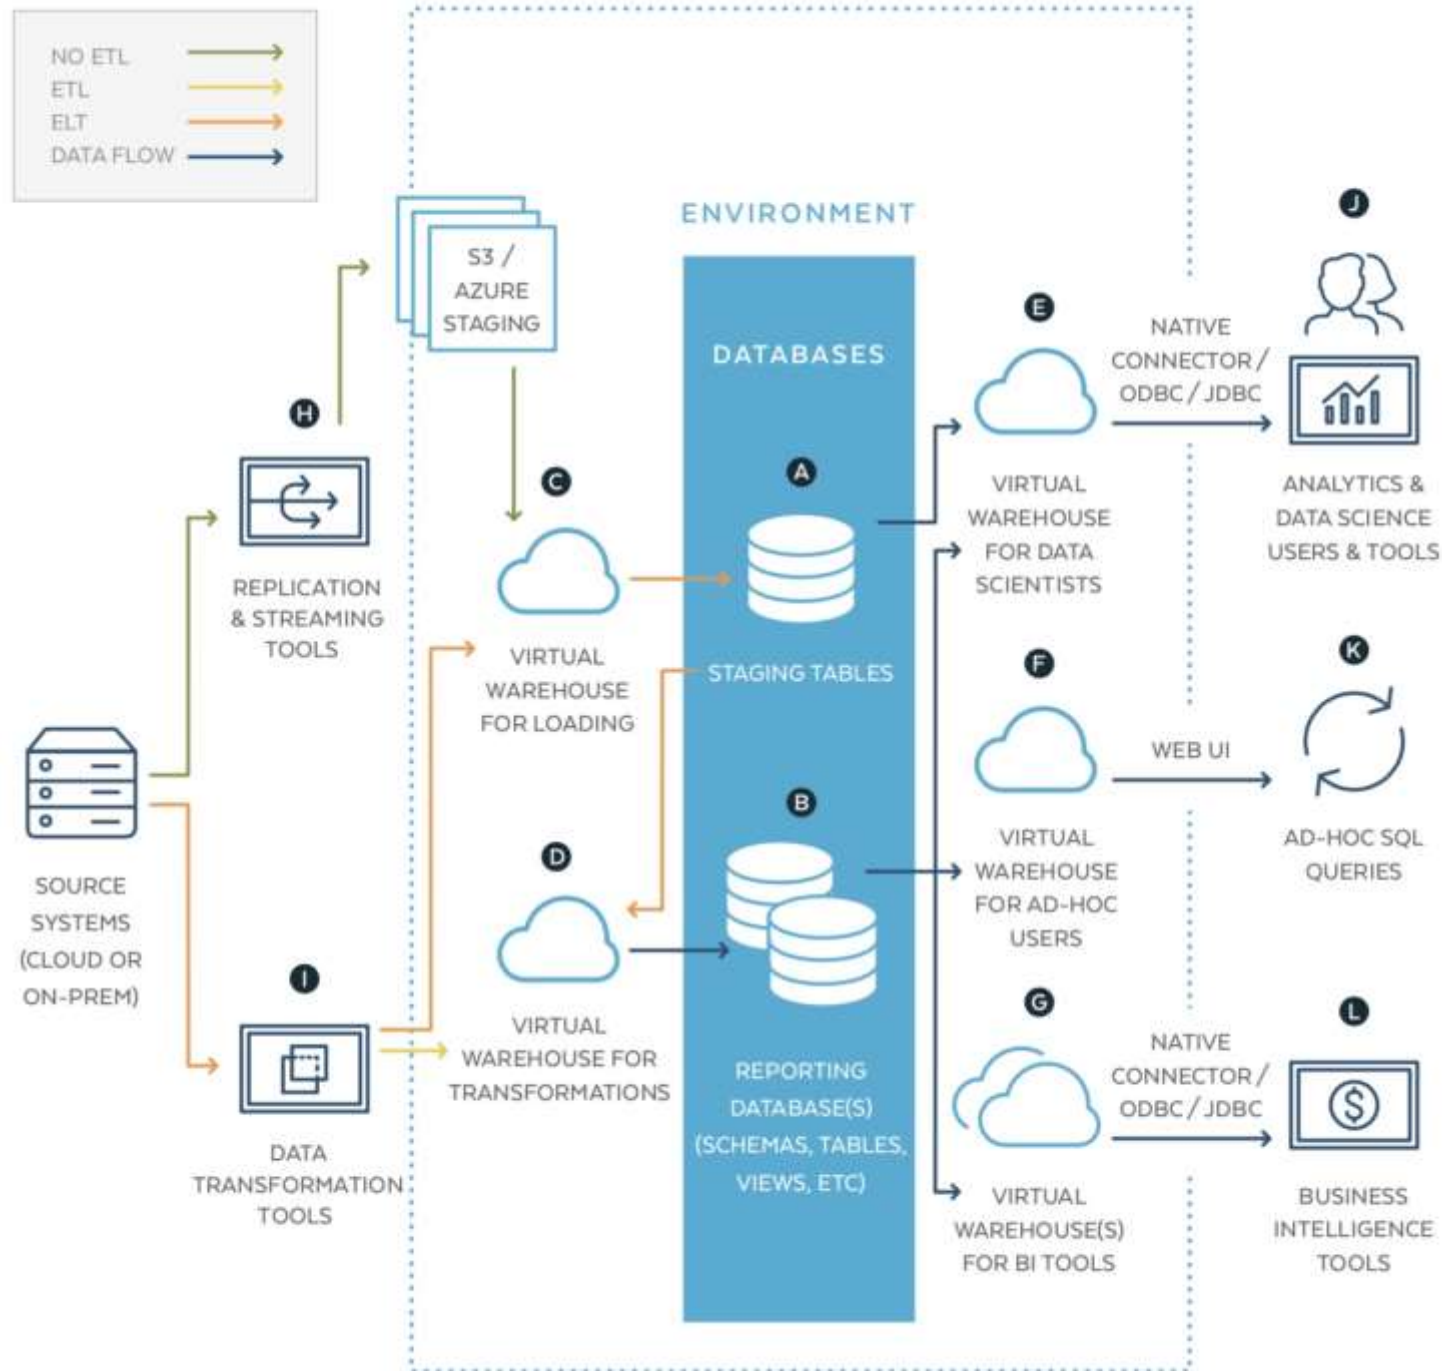


EXECUTING THE EVALUATION

Develop a Test Plan

- Identify appropriate level and scope of testing for each environment
- Automate testing so it's repeatable and provides results for identifying any issues
- Define, document and agree to acceptance criteria for the tests





EXECUTING THE EVALUATION

Load Initial Data Sets

- Create scalable but meaningful sample scenarios for processes
- Extract data from the existing data warehouse to S3/Azure Blob Storage, so that it can be loaded into the Cloud DW
- ..



EXECUTING THE EVALUATION

Take it to the Edge Cases

- > ..
- > ..that load data changes into the cloud Data Warehouse with appropriate process dependencies
- > Try a process that you cannot currently do on your existing implementation...
- >



EXECUTING THE EVALUATION

Yes, you need to test the evaluation

- Start test plan after initial data sets are loaded and data is being kept up-to-date
- Compare data between the existing data warehouse and Snowflake throughout the migration and resolve discrepancies
- Compare performance between the existing data warehouse and Snowflake share results with stakeholders



Functionality You *Must* Have

Don't mistake "familiarity" with a good product....



A TRUE EVALUATION

Run Existing Data Warehouse and Evaluation in Parallel

- Running the existing data warehouse and cloud data warehouse in parallel is relevant for testing the migration
- Minimize the amount of time both systems run in parallel, while spending sufficient time validating the capabilities



VALIDATING THE EVALUATION

Check Tools to the Cloud Data Warehouse

- Update tool connections to redirect to the CloudDW after validating data
- Copy existing solutions that point to the prior data warehouse and update them to point to the CloudDW
- Compare the output of the tools to validate that the results are the same



EVALUATION SUCCESS FACTORS

Identify Differences Between the existing data warehouse and CloudDW

Remember that familiar capabilities do not necessarily mean the product is not "better"

Resolve the Differences

Document functionality differences and shortcomings;
Establish a plan of action for reviewing issues and getting an updated status on resolving each issue

Communicate Evaluation Performance

Use the high-level outcomes captured while preparing for the evaluation to document the actual benefits that occurred



THE RESULTS MATTER

Document, Communicate, Monitor Differences
Between the prior data warehouse and CloudDW:

- Training Plan and Skills Assessment
- Upgrade tools and data driven processes
- Champion migration benefits / advantages



Discover the performance, concurrency, and simplicity of Snowflake

As easy as 1-2-3!

- 01 Visit Snowflake.net
- 02 Click “Try for Free”
- 03 Sign up & register

Snowflake is the only data warehouse built for the cloud. You can automatically scale compute up, out, or down—independent of storage. Plus, you have the power of a complete SQL database, with zero management, that can grow with you to support all of your data and all of your users. With Snowflake On Demand™, pay only for what you use.

Sign up and receive
\$400 worth of free
usage for 30 days!



Contact Info

Laura Ramsey
Snowflake Computing

Laura.Ramsey@Snowflake.com
@42LKR42

More info at
<http://snowflake.com>

Join my DBDEV FB Page:
<https://www.facebook.com/databasedevs>





THANK YOU

