





Big Data on AWS

Big Data Agility and Performance Delivered in the Cloud


Managed Services for Analytics



Retrospective
analysis and
reporting



Here-and-now
real-time processing
and dashboards



Predictions
to enable smart
applications

Elastic and highly scalable

+

No upfront capital expense

+

Only pay for what you use

+

Available on-demand

= the Cloud removes constraints

The Cloud Was Built for Big Data



Agility: Try more, fail fast, go big or start small, and process data at any scale



Scalability: Run jobs any time, without guessing capacity or limiting functionality



Broadest and Deepest Capabilities: Access 70+ managed Big Data services to address any workload



Low Cost: Pay only for the IT you use, when you use it

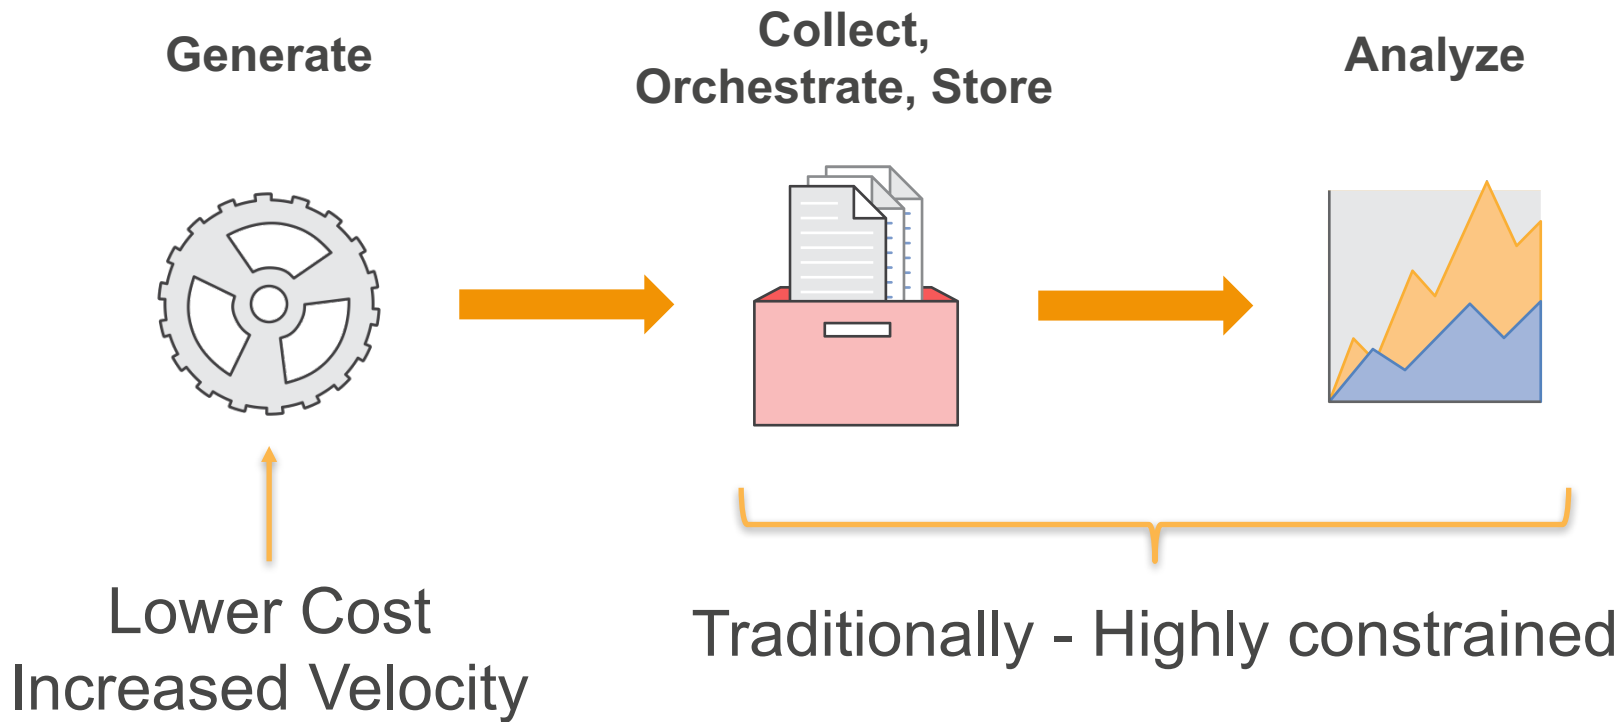


Get to Insights Faster: Focus on data science not the heavy undifferentiated lift of managing raw data



Data Migrations Made Easy: Move exabyte-scale data to the cloud quickly and cost-effectively

Common Big Data Flow



AWS Big Data Platform





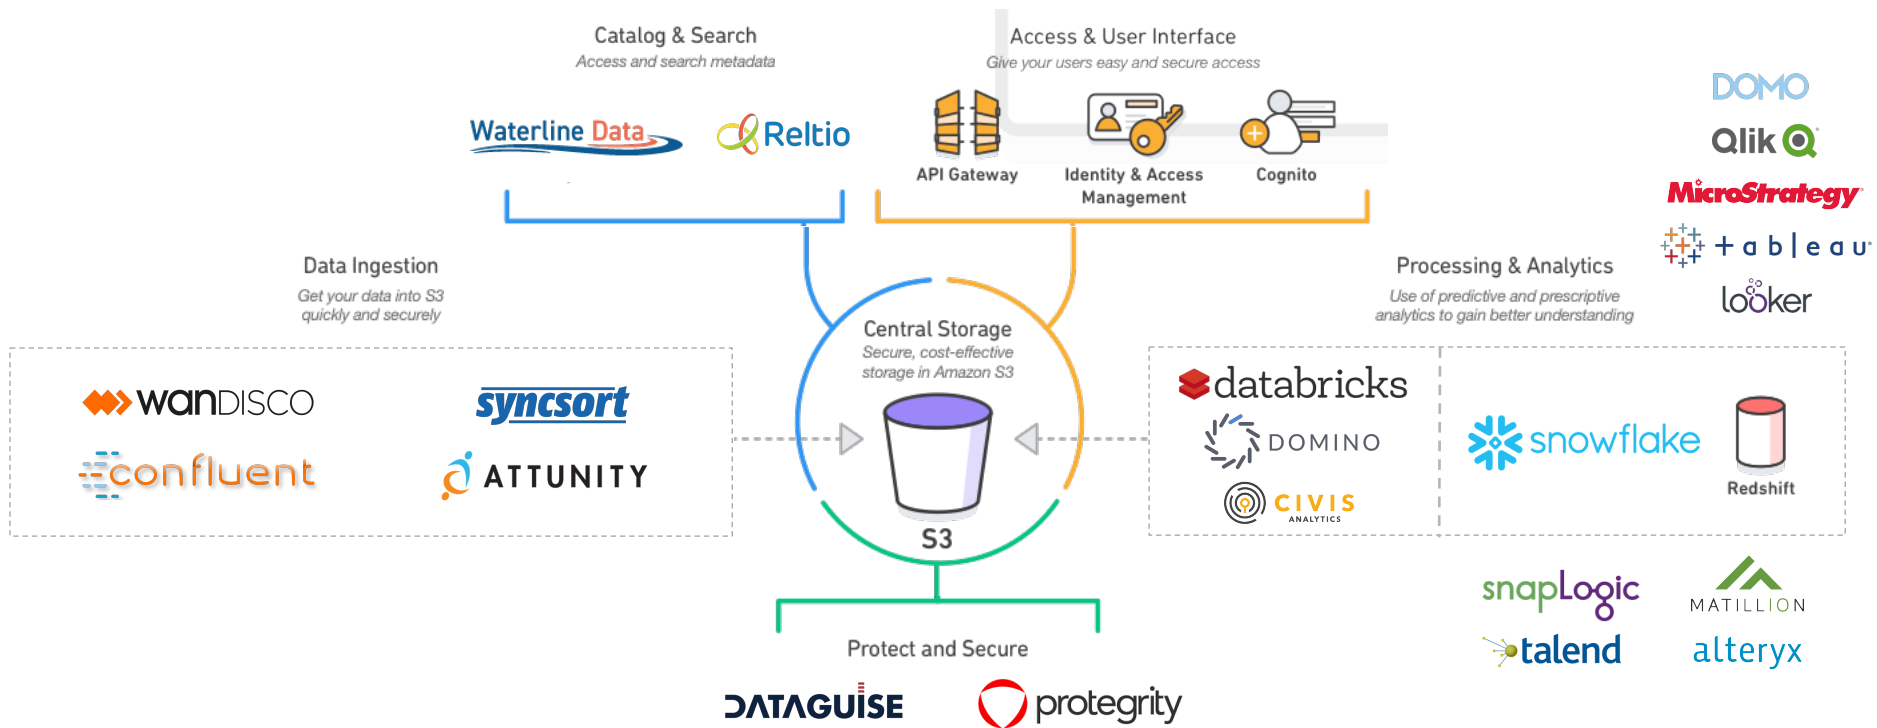
No one tool rules them all



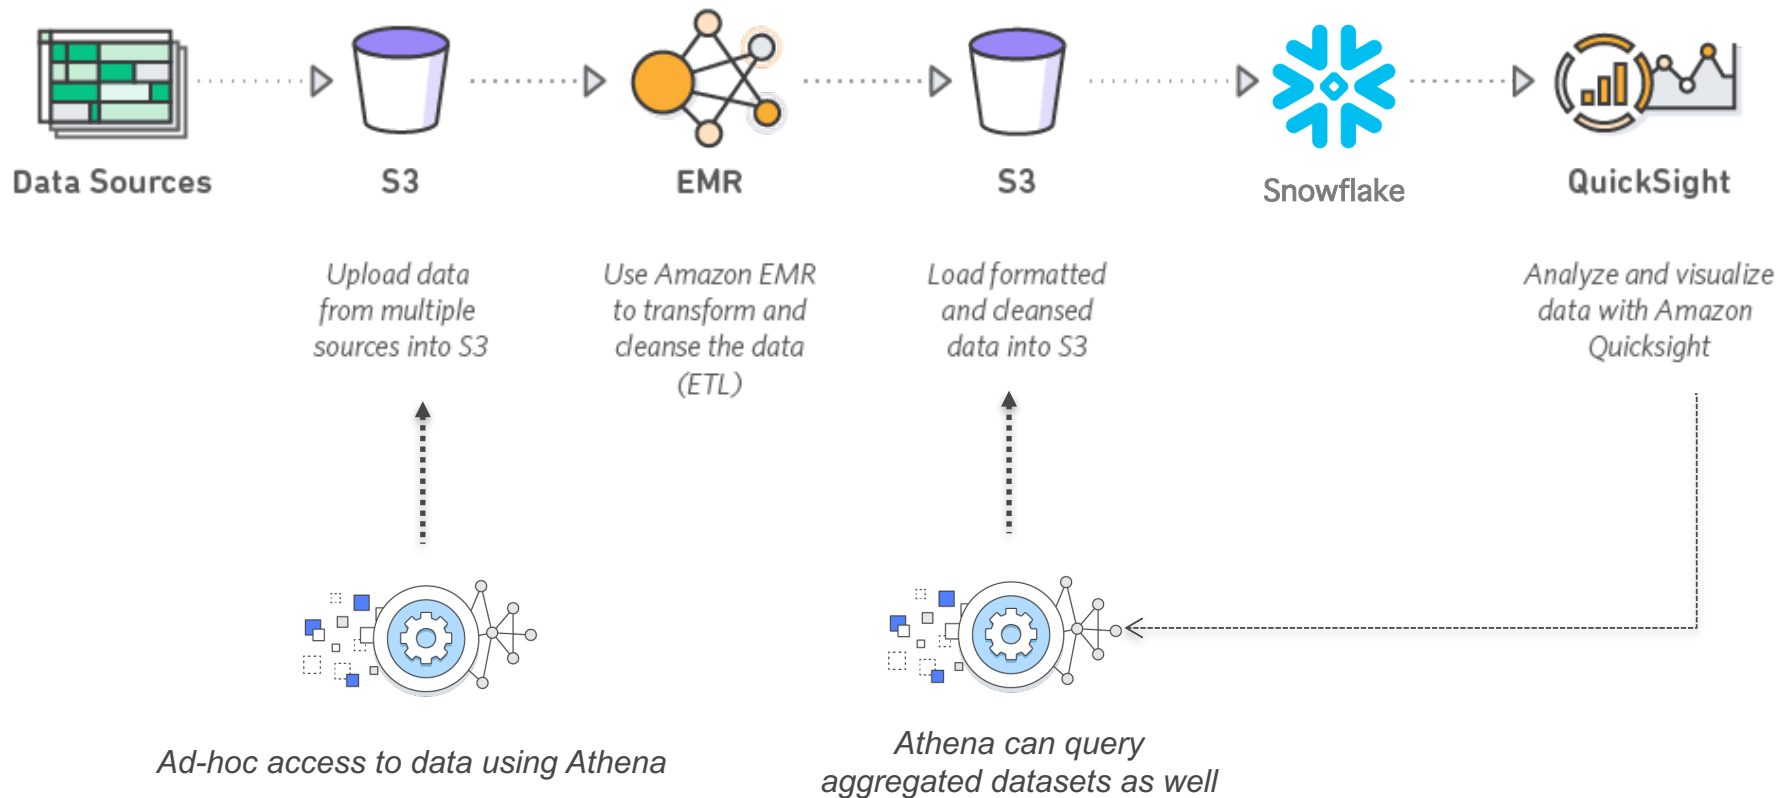
The AWS Approach

- **Flexible** - Use the best tool for the job
 - Data structure, latency, throughput, access patterns
- **Low Cost** - Big data \neq big cost
- **Scalable** – Data should be immutable (append-only)
 - Batch/speed/serving layer
- **Minimize Admin Overhead** - Leverage AWS managed services
 - No or very low admin
- **Be Agile** – Fail fast, test more, optimize Big Data at a lower cost

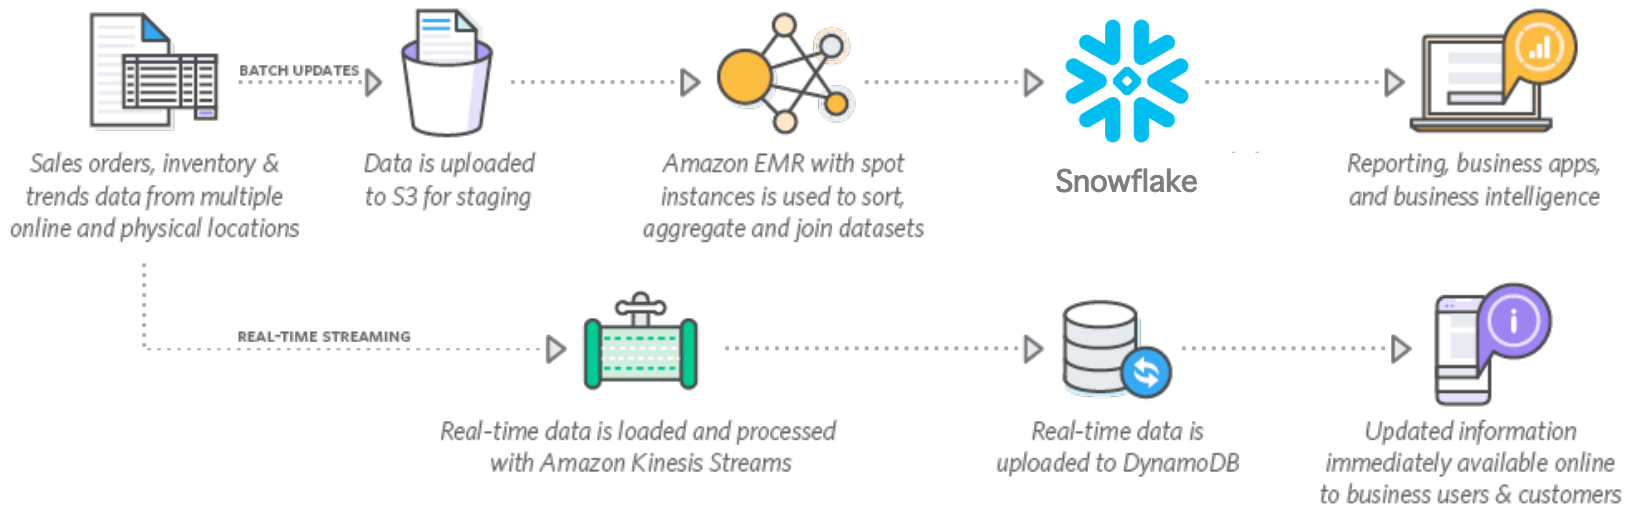
Building a Data Lake on AWS



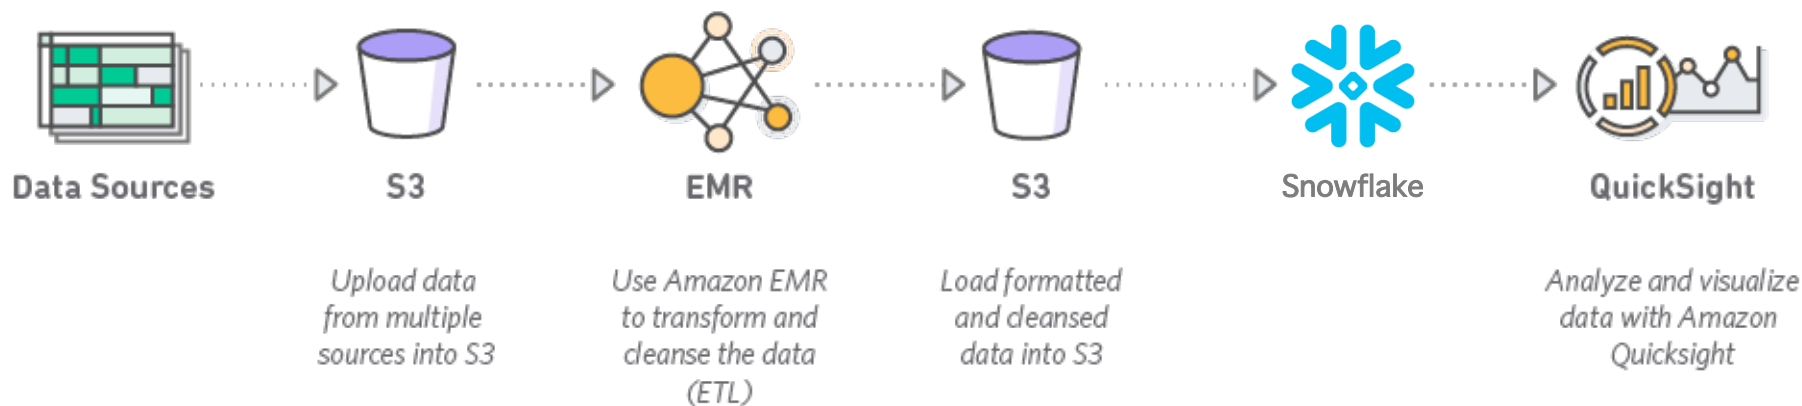
A Sample Batch Analytics Pipeline



On-demand Big Data Analytics



Data Warehousing



Summary



- Build **sophisticated Big Data applications cost-effectively** and support retrospective, real-time and predictive analysis
- You can **build incrementally**, scale automatically and add use cases as you go
- AWS delivers added benefits of **security** and **auditing** features to enable you to meet your stringent requirements
- Build **hybrid** applications that span across your datacenters and the AWS Cloud

Getting Started: Tutorials & Blog



Try AWS with 10-Minute Tutorials

10-Minute Tutorials are simple "Hello, World!" technical documents to help you get hands-on with AWS.



10-Minute Tutorial

Launch a Linux VM
using Amazon EC2



10-Minute Tutorial

Store and Retrieve a File
with Amazon S3



10-Minute Tutorial

Launch a WordPress Website
with Amazon EC2 and AWS Marketplace



10-Minute Tutorial
Launch a Web Application

with AWS Elastic Beanstalk



10-Minute Tutorial

Register a Domain Name
using Amazon EC2



10-Minute Tutorial

Store Multiple Files
to Amazon S3 using the AWS CLI



10-Minute Tutorial

Update a Web Application
with AWS Elastic Beanstalk



10-Minute Tutorial
Create and Query a NoSQL Table
with Amazon Dynamo DB

Subscribe to the **AWS Big Data Blog**: <http://blogs.aws.amazon.com/bigdata/>