



THE ROLE OF DATA WAREHOUSES IN BUSINESS ANALYTICS AND BUSINESS INTELLIGENCE



WHAT ARE DATA WAREHOUSES USED FOR?



The Big Data revolution was supposed to be a gold mine for organizations everywhere. Suddenly we'd have piles and piles of amazing, detailed information about our clients, our businesses, our industries - anything that could give us an edge - that could then be analyzed and turned into actionable insights to help us get ahead.



AS THE DATA PILES KEPT GROWING, THOUGH, THE PROBLEM QUICKLY BECAME CLEAR: WHERE ON EARTH DO YOU PUT ALL THAT DATA?

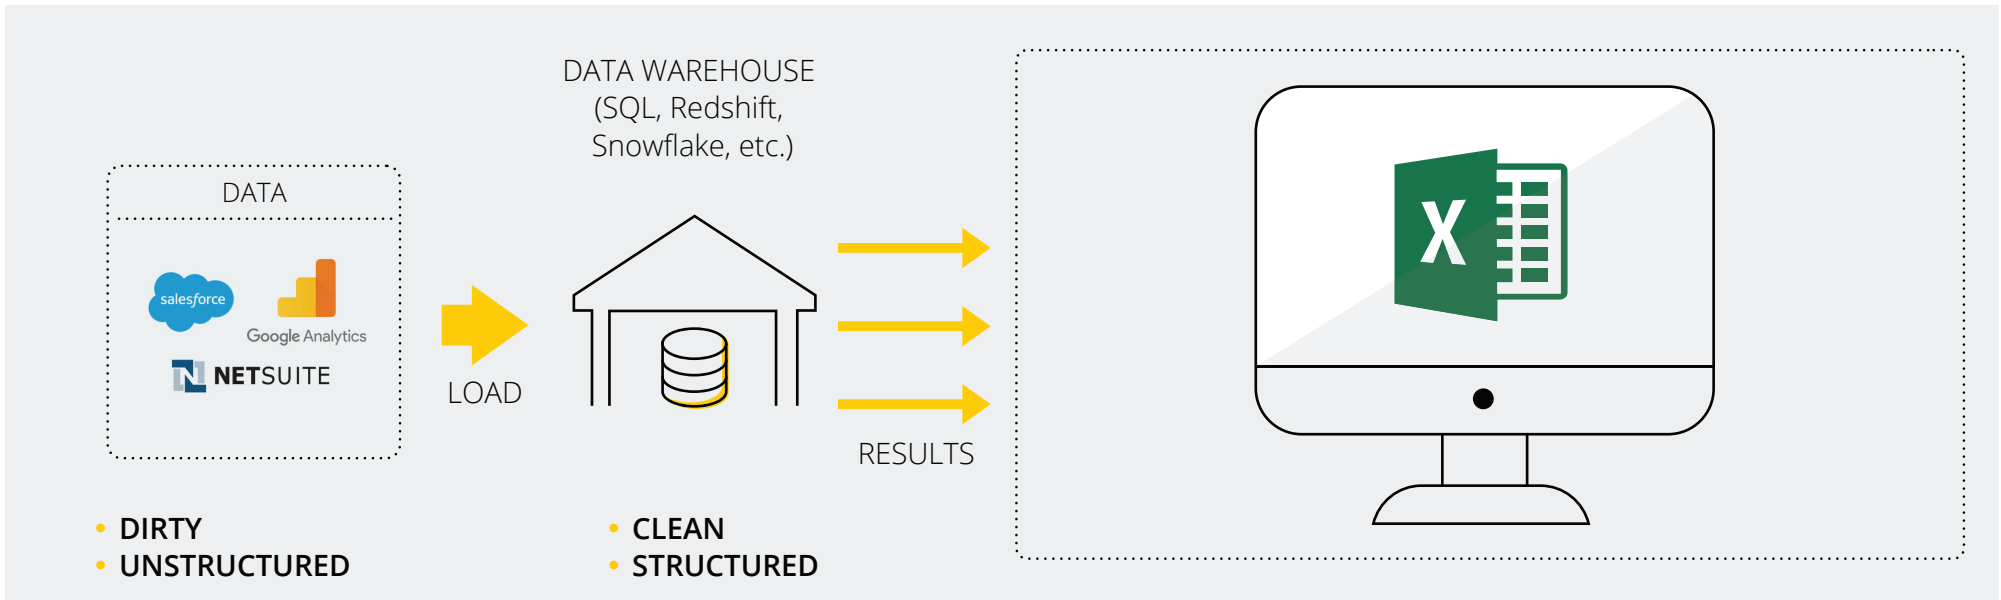
That's where data warehouses came in. These allow you to bring in all your data from all the myriad sources it comes from, stick it in one place, organize it in a clear and comprehensive way, and ultimately create a single version of truth. Major investments were made to make sure the data was stored in a way it could be pulled out quickly and conveniently with all the required aggregations prepared in advance. You could then use this as one source for Business Intelligence and Business Analytics.

... Aaaaand that is where more BI tools are needed. Your data warehouse might be the perfect storage system for all that data, but you still need something on top to extract it from there, and then a way to join the different sources, and then a tool to analyze the data - which means you need different tools, and people to write code and prepare reports that can allow you to work directly with the data stored in your data warehouse. Something that will allow users to dive right in, find what they need, analyze and draw out the key insights to make better business decisions.

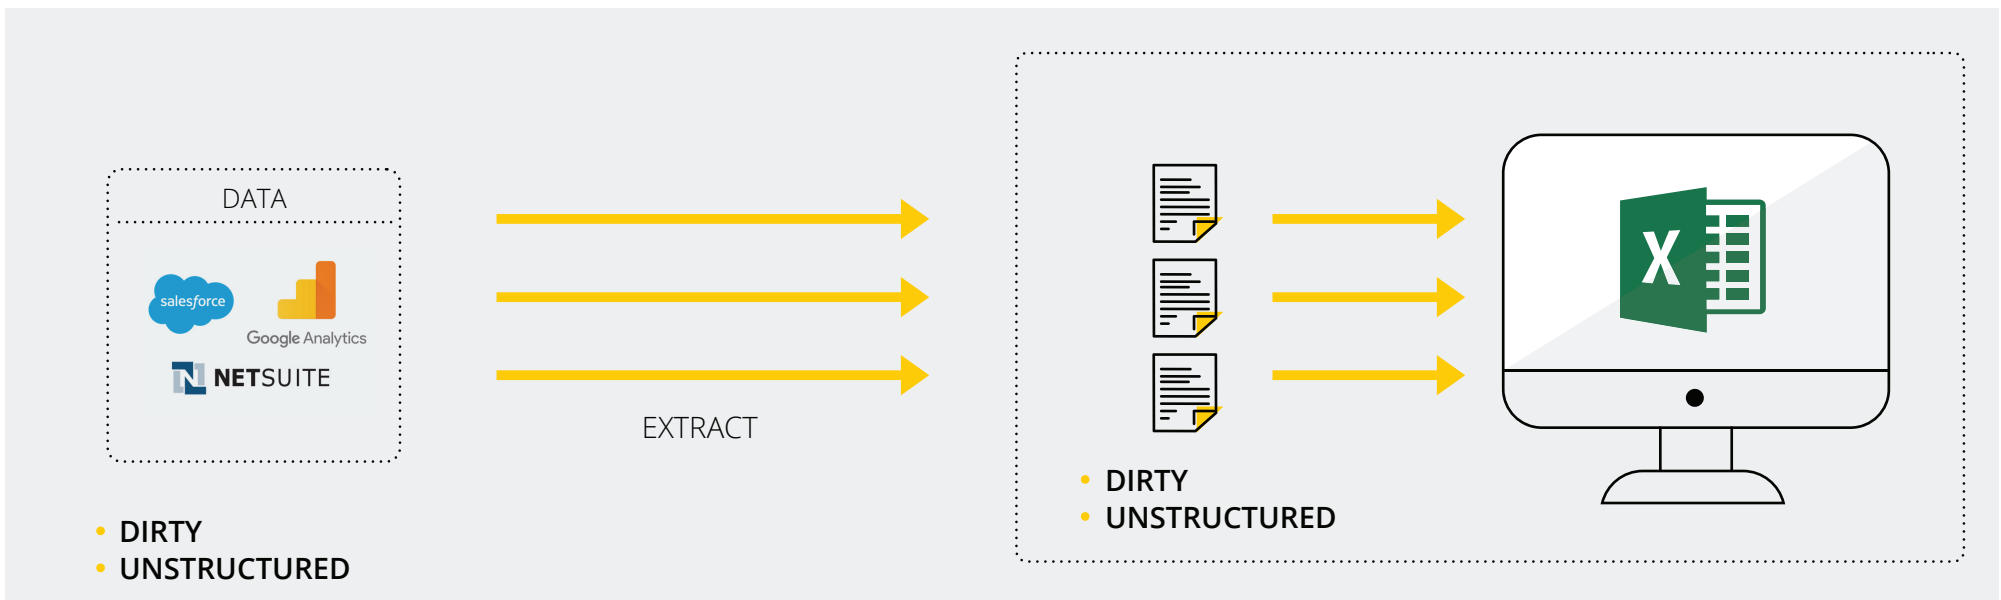
For example, there are solutions out there that allow you to extract the data you need, manipulate and analyze it quickly and conveniently, and refresh this swiftly whenever your data is updated. With this type of technology, you may even be able to cut straight to the original applications without having to load everything into your data warehouse first. But more on that later.

How Data Warehouses Work

DATA WAREHOUSE



NO DATA WAREHOUSE



In the past, businesses without a data warehouse were stuck with masses of highly detailed, but messy and unstructured data, from which end users on the business side of the organization struggled to get anything useful out of. A data warehouse takes all that scattered data and lays it out nice and tidy in a huge relational database, through a process of Extract-Transform-Load (ETL). This takes a huge load off transactional systems, as well as improving data quality and helping to prepare it for analysis. This also makes sure that when extracting or adding data the operational systems are not being affected.

The trouble with this becomes clear when you try to actually use that data for Business Intelligence (BI). No analytics engine is strong enough to run queries that sift through every fragment of data contained in your database. To do that, you need a phenomenally powerful hardware setup, which would cost a fortune!

To get around this, organizations started using “data marts”. These work by giving different groups of people access to one subsection of the whole data warehouse. While this reduces the pressure on an individual’s machine when they need to run a query, it also means no one person can ever get the full picture from the data the company collects. Hardly an ideal outcome.

Traditional BI didn’t have a way around this, though. It couldn’t tap into the raw data in the data warehouse, or pick and choose what it wanted to see within the database. As a result, the BI solutions of the past used either data marts or OLAP cubes to connect to the data. The latter cuts down query times by pre-summarizing certain elements of the data before it starts searching through the database.

BOTH OPTIONS ARE SPEEDIER THAN WORKING WITH COMPLETE, GIANT DATASETS, BUT THEY MEAN YOU LOSE EITHER SCOPE OR DETAIL IN YOUR DATA, AS WELL AS DECENTRALIZING IT... WHICH PARTLY DEFEATS THE POINT OF HAVING A COMPREHENSIVE DATA WAREHOUSE IN THE FIRST PLACE.

What Can't a Data Warehouse Do?

DATA WAREHOUSES OFFER AN EXCELLENT WAY TO CLEAN UP AND STORE YOUR DATA IN A WAY THAT CAN BE EASILY SEARCHED AND ORGANIZED. THAT'S WHAT THEY ARE FOR.

The important thing to understand, though, is that this is all they are for. They don't have any analytical function, so you need that on top. The fact that data warehouses are so good at tidying up your data can also be a drawback. This is because they are relational databases, which makes them highly restrictive in terms of the kind of data you can store and how you store it.

You only have a certain number of columns and ways to sort and identify each item, so anything that doesn't fit neatly into this poses a challenge - for example, photo or video content, language analysis, and so on. It also means that data warehouses only store past/historical data, also referred to as analytical data. You can't store current real time (i.e. transactional) data.

This limits the kind of insights you can get out of the data stored in your warehouse, even with an analytics engine attached. Depending on what sources you use, getting a complete and accurate picture might mean drawing data directly from other applications or types of databases into your BI platform.

Then there's the issue of getting enough support from IT. If you're moving with the times and embracing new ways of engaging with customers or collecting information about your business and processes, that means you either need to add new data streams and types to your warehouse constantly, or you'll need to work with multiple systems.

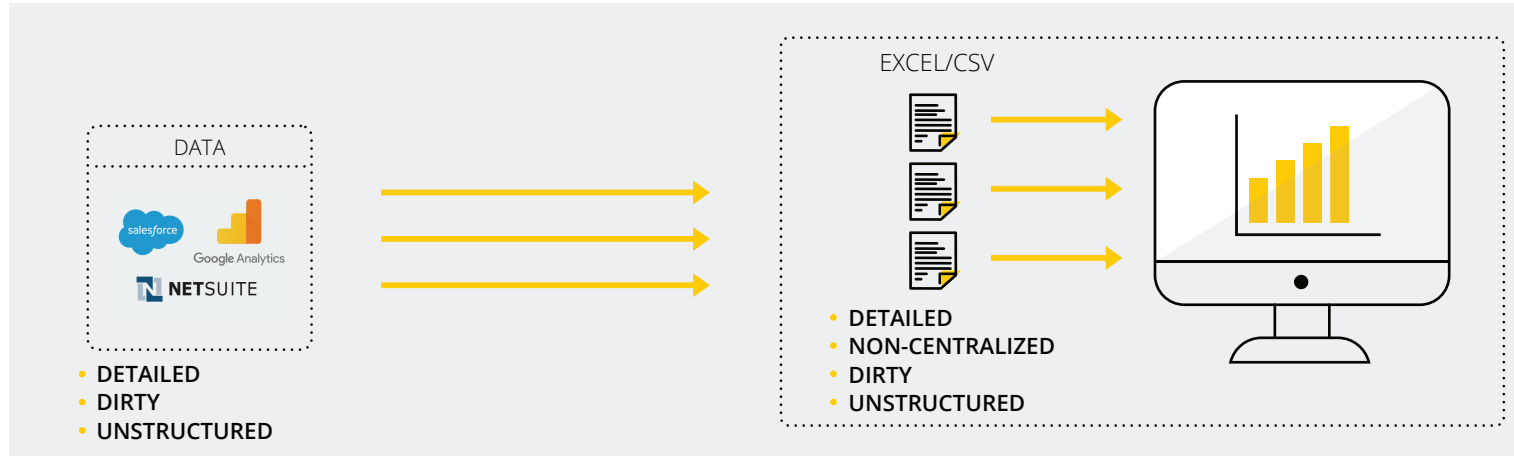
The former can feel frustrating if you've already invested heavily in building a data warehouse, while the latter is far from a simple task - you'll almost certainly need plenty of help from your IT department to keep on top of it. If you do, and this creates a delay in accessing and analyzing all the data available to you, you could find that your data warehouse ends up slowing you down instead of speeding you up. This situation simply means that you have to keep investing in your data warehouse to make room for more data and to be able to load data from new data sources that did not exist originally when the data warehouse was built. A good example is a service provider that stores billing information regularly. When a new service is added, the data from this new service is not uploaded to the data warehouse until it is configured to do so by programmers. Until then, the data can not be analyzed, leaving management in the dark about adoption and success.

Does Everyone Need a Data Warehouse?

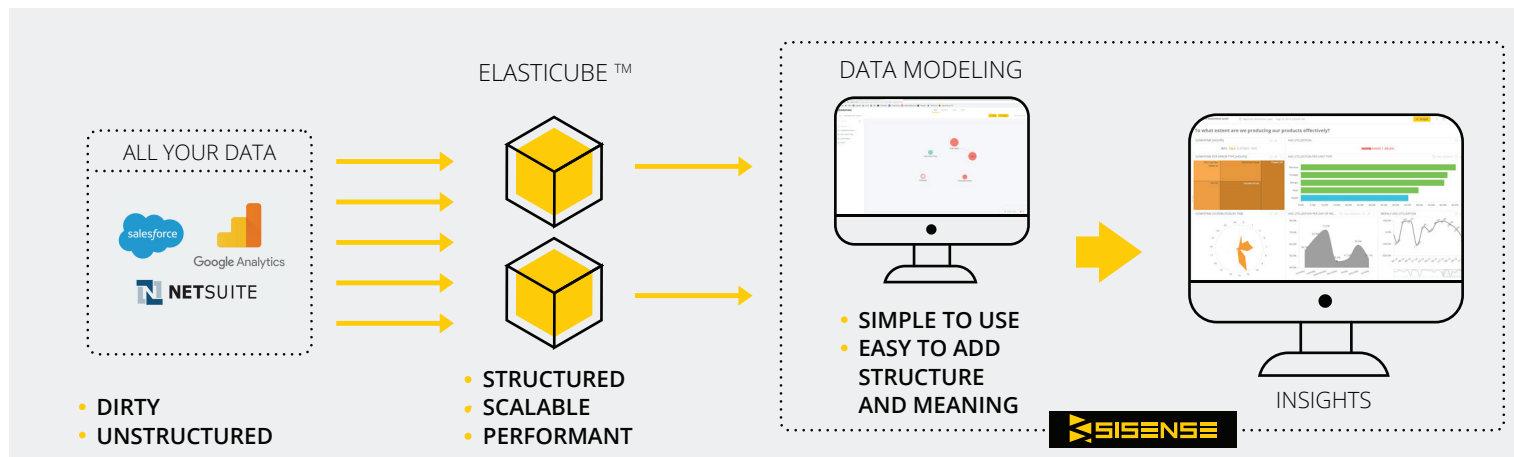
In the last section, we mentioned that your data warehouse's limitations could mean you need to use other data streams alongside the warehouse for analysis. However, if you have a sophisticated enough BI system that you can connect to those original data sources directly, and can guarantee your access to clean and complete data... well, you might not need a data warehouse at all.

Take Sisense, for example. Their solution to this issue is to build ElastiCubes - unique, high-performance analytics databases with super-fast data stores. ElastiCubes use an in-memory columnar database approach, meaning that a) rather than having to load all the data in the database into the BI platform for analysis, it whips through and loads just the bits you need, and b) it optimizes use of your computer's disk, RAM and CPU, so that you can use a single commodity server to crunch terabytes of data on a standard machine, even if there are a high number of concurrent users.

TRADITIONAL BI (WITHOUT DW)



SISENSE WITHOUT DATA WAREHOUSE

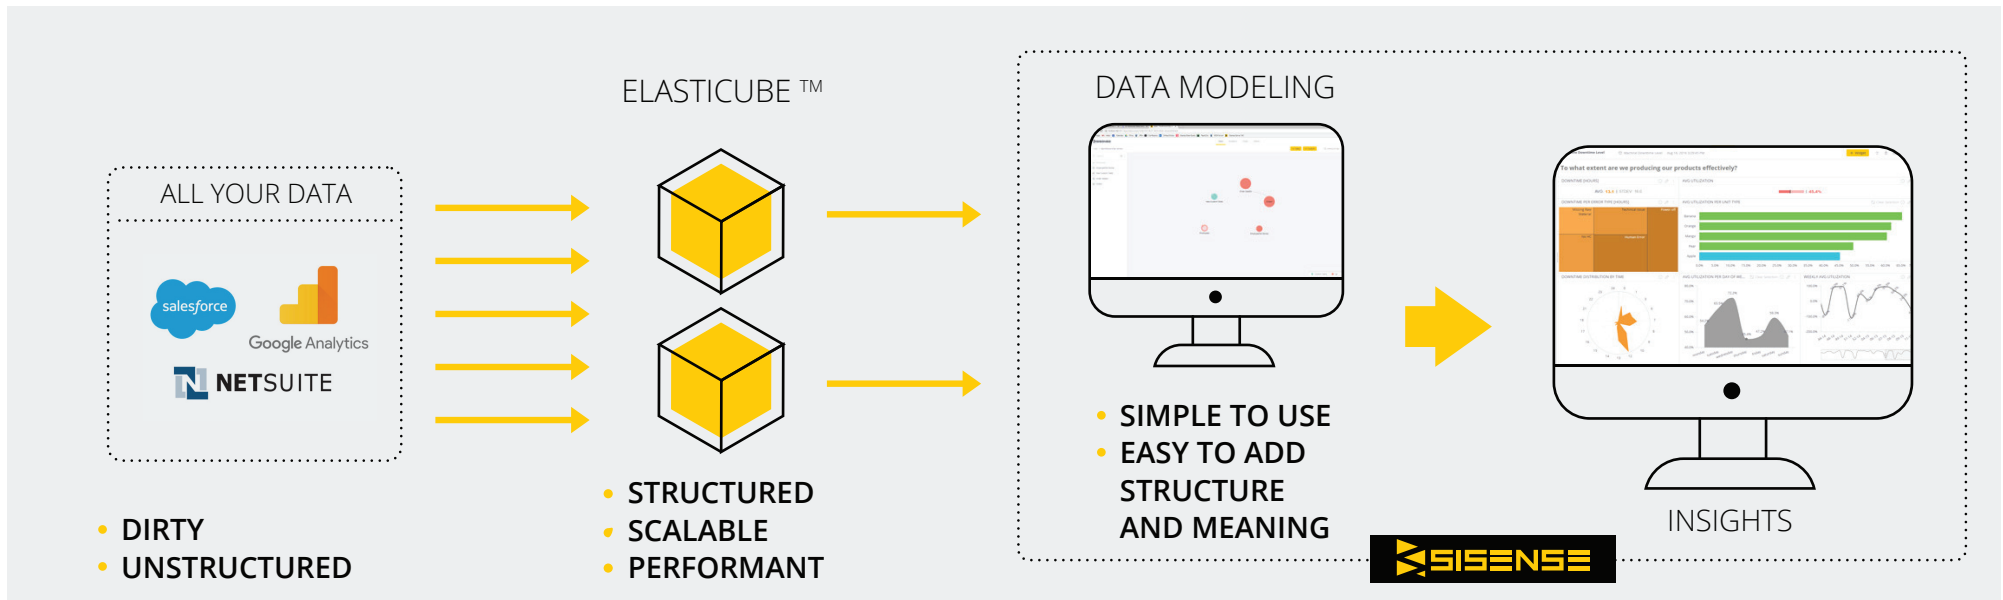


It also cleans, centralizes and structures this detailed data for you, performs ETL functions within the ElastiCube server, and presents the results in an intuitive dashboard so that it's optimized for use by business users of BI. This dashboard can easily be shared across your company.

As a result, you can access all the data you have for any given analytics task, in a neat and tidy format, without having to invest heavily in hardware or lean on IT for support. If you don't yet have a data warehouse, you may find that this system makes one redundant.

SISENSE

WITHOUT DATA WAREHOUSE



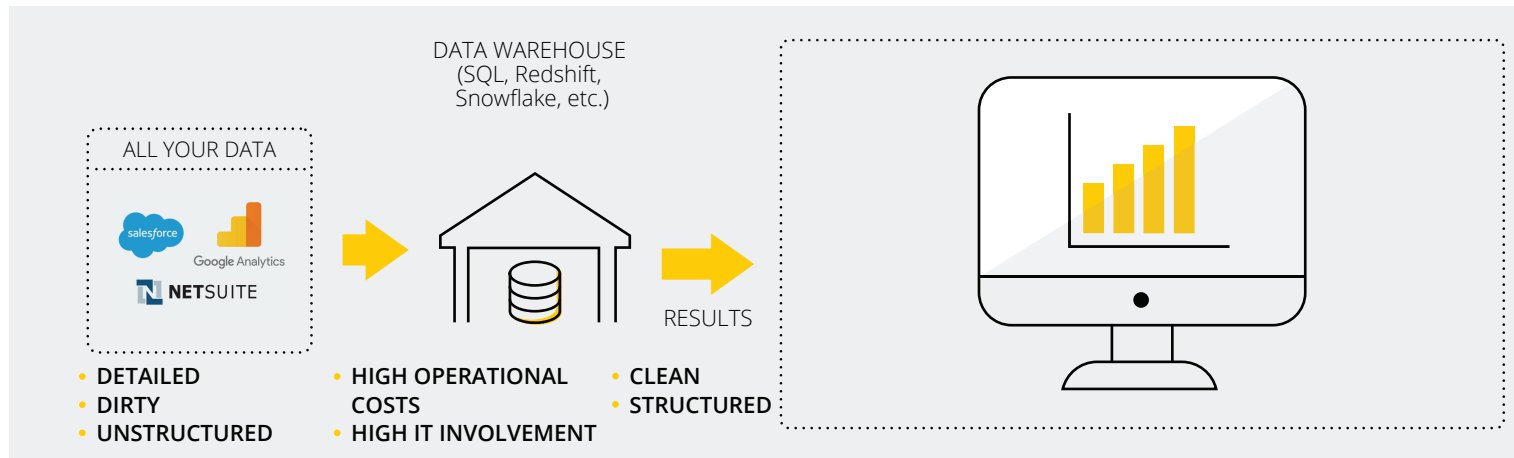
But... What if All Your Data is Already in a Warehouse?

If you already have a data warehouse, though, your efforts haven't been in vain. In fact, it can even be a good thing.

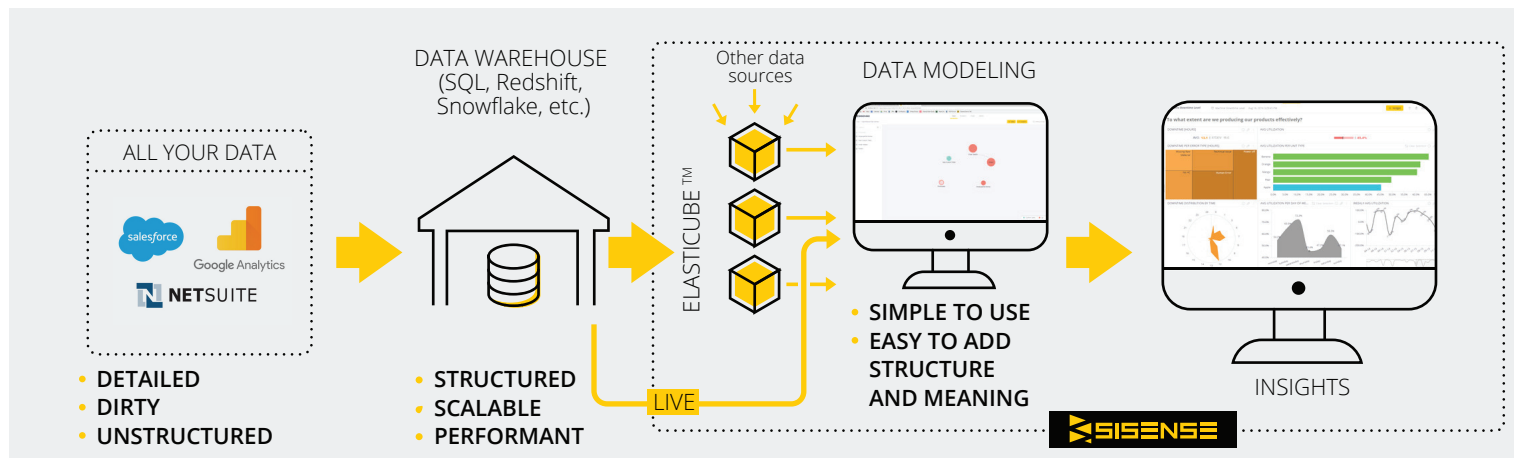
If you've already committed to a data warehouse, there's no reason you can't continue to connect this to a BI system for analysis. In fact, the right BI solution can actually help you to get more out of your existing data warehouse, to maximize your investment.

Again, to use the Sisense example, in this setup you can connect ElastiCubes to the raw data in your data warehouse, allowing you to retain the granular detail of this clean, centralized, structured data. Compare this to traditional systems which use data marts or OLAP cubes to connect to this raw data, summarizing it and decentralizing it in the process. New, sophisticated BI technologies aren't here to render your database obsolete (unless you want them to). More commonly, they help you get the most out of the system you already have.

BI WITH DATA WAREHOUSE



SISENSE WITH DATA WAREHOUSE

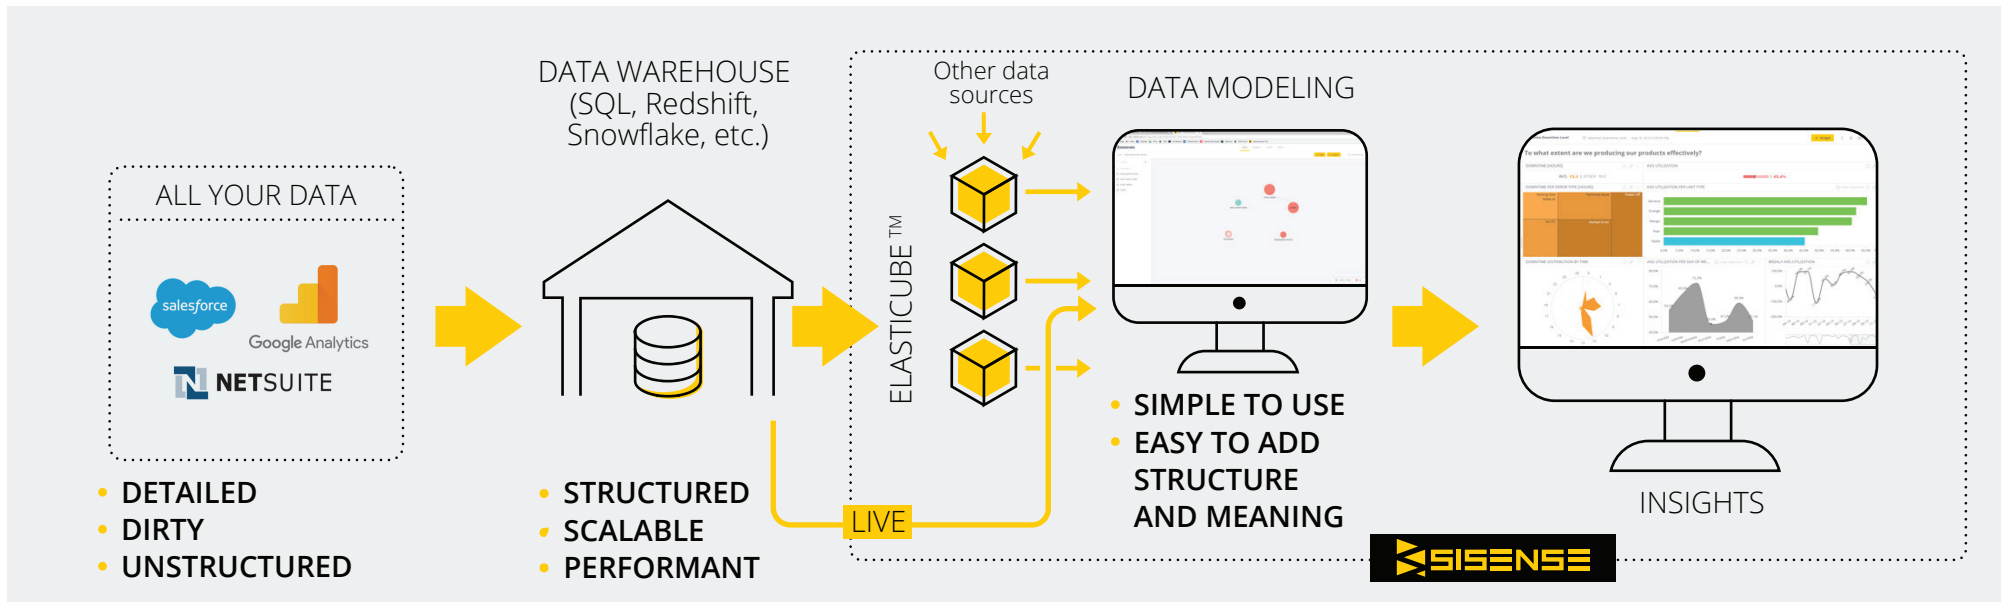


It does all this while future-proofing your data strategy. You don't need to establish additional ways to import new or difficult data streams into your database because you'll be able to connect to them directly.

You don't need to worry that your store of data is growing all the time, because it's not going to bottleneck when you try to get it from your warehouse and into your analytics tools. You can scale.

SISENSE

WITH DATA WAREHOUSE



Final Thoughts: Getting Smart With Your Investment

Until recently, the first step in any organization's BI efforts was to build infrastructure for collecting and storing data. Typically this meant investing in a data warehouse, which keeps growing as you continue to add data. The data stored in the data warehouse would then feed into the BI platform for analysis. Investments in data warehouses and the tools to extract the data and analyze it were massive.

In today's BI world, you can still do that - but it's not the only way. You can also make direct, even live, connections to your data sources, without necessarily analyzing only the data that was stored in a data warehouse.

This means that a data warehouse is no longer necessarily central to your BI process, but that doesn't mean it's obsolete. The best practice is to use your BI platform to mash up data from data warehouses and direct connections in one place.

Using a data warehouse with a BI platform like Sisense gives you the best of both worlds. You continue to use your existing setup without incurring new hassles or paying out again when you want to add new data streams, plus you actually make the data in the warehouse more accessible, useful and ultimately more valuable than ever before.



Explore Data in Sample Dashboards

See interactive BI dashboard reports in action on sample data.

[Interact with Live Dashboards](#)



Get a Real Feel for Sisense

Interact with a live Sisense BI dashboard - drill down, filter and customize your data.

[Start Your Free Trial](#)



Take the Ultimate Test

Go from data to dashboard with your own data in just 90 minutes.

[Schedule a Custom POC](#)