

Why have I never achieved the 360 degree view of my customer?

Why have I never achieved the 360 degree view of my customer?

One of the most common goals for enterprises today is to achieve a single view of their customer. The potential benefits are both far-reaching and significant, as different departments including sales, marketing, customer success and finance are able to use this holistic view to deliver higher levels of customer satisfaction, build deeper relationships with customers and predict who their future customers will be.

Despite unilateral agreement that the single customer view is highly desirable, most enterprises have failed to achieve it. Why is this? In many cases, it isn't for lack of trying. There are a vast number of products and services available to buy today which promise to deliver the single view of the customer. However, the majority of them can only enable part of the picture, and require a huge amount of time and effort from internal teams to do so.

In this paper we will examine why the single customer view has proved to be so elusive, and how to overcome the common barriers. We'll also outline not only how companies can finally achieve a 360 degree view of their customer, but also how these principles can be extended to deliver even greater value to the business.

What is the single view of the customer?

For enterprise customers, the single view of the customer is the profile that is created by the aggregation of all the records you hold across your systems that relate to one customer. This could be anything – including contact details, transaction history, social media interactions, email communications and meetings. These records then need to be collected, cleaned and enriched before combining to create one Golden Record, or single source of truth.

Common challenges with creating a single view

Many of the issues enterprises encounter when attempting to create a single customer view stem from preconditioned notions of where to begin.

Traditional master data management dictates that you must model your data upfront, which means deciding what a customer should look like and how it relates to other domains before you have even begun to discover what the data can reveal about your customers. If you think about it, it's completely counter-intuitive. You're embarking on this project because you would like to uncover hidden characteristics and insights about your customers. Yet you're being asked to determine what this is going to look like before the project gets underway. By creating this model upfront, you've already limited the potential to learn and act upon the intelligence you'll gain.

There are other complications too. The view of your customer record on a metadata level can be composed from many different sources. Blending it and understanding that two records are the same person is not easy. So-called "single view" technologies and legacy master data management systems will struggle to do this.

Dirty data – incomplete, stale or inaccurate records in this case – only compounds the issue and makes it more likely that records which do not necessarily relate to the same customer will be integrated and presented as an aggregated view. Without full tracking and lineage of where the data came from and how it was modified over time, you cannot have confidence and trust in the view you eventually get.

Perhaps the biggest challenge of all is that single customer view projects can very quickly spiral to become more complex than originally predicted. The customer domain does not exist in isolation, and is connected to other domains such as products, purchases, and so on. If these associated data domains have not also been sufficiently cleaned, integrated and enriched then at best they will compromise your complete customer view, and at worst render it invalid.

Despite the various difficulties outlined above, the fact remains that every organisation on the planet could benefit from knowing more about its current and future customers. The less you know about them, the higher the chance that you will fail to attract and retain customers, and the greater the possibility that your competition will fill the void you have left.

When good is good enough

As valuable as a single customer view can be, you will almost certainly never be able to point to a single Golden Record and rely on it as a complete and wholly accurate representation of the customer. But that's okay. Even if you had every single record from the past five years, associated and available as one single view, would you actually use all of that data? Do you need to know that four years ago they submitted two support tickets which took over 48 hours to respond to? If it's still taking over 48 hours then yes, you would probably want to know that, but if the issue has been resolved then maybe you don't need to. Perhaps, for now, it's good enough to know that over the past two years they haven't increased the number of licenses they have or that they haven't actually activated part of their subscription. Maybe that is the prime opportunity to engage account managers and customer success teams to help them get the best value out of the product and potentially bring more users onboard. The fact is that you don't have to know everything there is to know about your customer before you can use the insights delivered by the data you do have.

A single view of...anything?

As alluded to above, trying to model what your single customer view looks like before you embark on a master data management project will only set you up for failure. Data, just like your customers, is a living, evolving entity. In a way, you will always be playing catch-up with your data and the customers it relates to. If you adhere to the model upfront approach, with its strict constraints and limitations, you will miss a lot of the nuances and hidden insights contained within the data. Data should be allowed to connect and relate in a natural rather than being forced. Truly modern Master Data Management platforms enable this because they are based on a Graph database, which acknowledges that the relationships between the data are as important as the data itself. Think of it as storing data in columns and tables in Excel versus jotting down your ideas on a whiteboard – it's far easier to visualise and make the connections when the data isn't restricted by predefined fields. As a NoSQL, schemaless database, Graph is the critical element of modern approaches to mastering data.

Of course, this applies to every domain, not just that of the customer. By broadening your focus from the single view of the customer to the single view of the product, or transaction, or location, you will exponentially increase the value that data delivers to your business. As well as gaining a better understanding of your customers.

A great example is the organisation that wanted to know more about their customer support department. They wanted to know what makes a happy customer. On the surface, this is the perfect use case for the single customer view. But in order to understand what made the customer happy, they also needed the flexibility to pivot around other anchor points. A customer cannot tell you what makes them happy at the aggregate level, so you need to be able to uncover what is happening in relation to other data objects – like dates, transactions and departments.

Unless you have built the right data foundation to enable this, trying to shift the lens to focus on other angles is going to be virtually impossible. This isn't due to a lack of foresight on behalf of the enterprise, rather that the technology simply wasn't available until recently.

In the old world, before Graph-enabled Zero Modelling, there were always limitations on the depth at which we could explore the data that is connected to a record. The single view of the customer cannot be achieved if we only think about records that are directly connected, because so many insights come from the connections that are indirect.

Predicting the pulse of tomorrow

Mastering any data domain is challenging for a number of reasons. Customer data in particular will change on an ongoing basis simply because it relates to human beings (who are notoriously fickle and unpredictable!). Instead of working against the changeable nature of data, the trick is to embrace it and build a data fabric that can adapt to those changes in an automated and scalable way.

Those that adopt this approach have an unprecedented opportunity to capture the hidden signals buried within the data and learn what customers want and how they want to receive it. Armed with this, you have a direct line into the pulse of the market and be able to answer it with a speed and scale that your competitors cannot. Once you tap into this customer pulse, you can start to use all of your data to drive more focused initiatives, serve customers better, drive cost efficiencies and so much more. Perhaps most significantly, the single customer view will help you to predict who your future customers will be, based upon the profile and behaviours of your happiest customers today. Armed with that information, the opportunities to grow your business will become abundantly clear.

