

**YouGov's** mission is built on the premise that the more people are able to participate in the decisions made by the institutions that serve them, the better these decisions will be. Institutions come to YouGov to find out the beliefs, behaviours and brand opinions of the people they serve via online surveys managed by YouGov. But to enable people to participate YouGov needed to make sure it was easy for them to take surveys. YouGov saw incomplete surveys by panellists when the speed at which survey questions were served up from the database slowed, due to inadequate storage performance. Today with Pure Storage installed, surveys are 80% faster at the storage level and abandon rates have fallen.

**BUSINESS TRANSFORMATION**

Faster response for panellists, quicker insight into survey data

**GEO**

United Kingdom

**INDUSTRY**

Market Research

**“Pure Storage’s solution allows us to provide a more consistent service to the analysis staff, which previously could be very disruptive and cause user frustration. We are a lot more consistent and predictable now; I haven’t had a problem since we went live in January. Previously, there was a download performance issue every week”**

*Nick Carter, Head of Infrastructure and System Operations, YouGov*

Small margins count for YouGov when it’s surveying voters, consumers and decision-makers – and a lapse of tiny fraction of a second can make the difference between a successfully completed survey and one that’s abandoned halfway through out of frustration. For Nick Carter, Head of Infrastructure and System Operations at YouGov, the recent decision to adopt Pure Storage’s FA405 is making a big difference.

**A BETTER SOLUTION NEEDED – AND ITS IMPACT ON GROWTH AND BUSINESS**

YouGov’s existing SAN was doing a reasonable job, but that was just the problem – the team had little scope for improvement. The company’s infrastructure is based around two rather large, terabyte-scale databases, one of which collects the responses to its surveys, the other managing its panel’s profile information.

To give an idea of the scale of YouGov’s operation, take this example from the 2015 UK General Election: YouGov polled 30,000 people at seven-day intervals to understand their voting intentions and track the changes in sentiment.

Such a heavy workload places huge importance on both reliability and performance; get it right, and the panel of 30,000 voters would continue to contribute happily. But, on top of that, the collected data would need to be processed and interpreted at speed to generate the timely, accurate insight that YouGov is famed for.

Carter knew that the quicker response times were when rendering survey questions, the better. Panellists in YouGov surveys are asked to answer a series of questions on a web page, and even a fraction of a second’s delay can affect how happy they are to complete surveys, and how many surveys they are willing to undertake. With the existing HD SAN running out of space, Carter was faced with a decision: buy more SAN space or buy an alternative.

Carter started a conversation with Pure Storage, and ultimately decided on a solution that would not replace YouGov’s existing SAN but complement it. Through implementing the FlashArray 405, YouGov has been able to move a lot of the more intensive workloads off the SAN, resulting in high IOPS and low latency on Pure Storage. The business’s main production store is MongoDB – currently a 1.3TB database with two replicates. A further two Postgres databases run to 1TB and 150GB.

Carter said that the team has “definitely seen a speed improvement, which has been very noticeable. We’ve also moved MS SQL over and had great performance”.

**COMPANY:**

YouGov  
<https://yougov.co.uk>

**USE CASE:**

- Database – MongoDB

**CHALLENGES:**

- Ensure fast system response times for panellists
- Manage storage requirements of two large databases (terabyte scale)
- Faster turnaround of survey results following election debates

**IT TRANSFORMATION:**

- Pure Storage FlashArray achieving 3:3:1 data reduction
- 658GB in shared space
- 3.19TB, using 2.31TB at the moment, 1.67TB is unique data

“One of the biggest positives from our perspective has been the account management team we’ve been working with. They’ve been exceptional, some of the best people I’ve worked with from a client management perspective. The Pure team are great guys, which makes them a pleasure to do business with”

*Nick Carter, Head of Infrastructure and System Operations, YouGov*

**KEEPING PANELLISTS ENGAGED AND HAPPY**

YouGov has more than 300,000 active panellists in the UK, and up to three million worldwide. The company is constantly collecting interviews, equating to more than 20 million responses a year. All of this data needs to be stored and retrieved as quickly as possible, to ensure the panellist’s experience is as smooth and uninterrupted as possible.

In a business that relies on a quick turnaround of data, and ensuring end user engagement, mere milliseconds can make a world of difference. Since implementing Pure Storage FlashArray, panellists can now move from one page to another faster than ever, with an improvement of more than 200ms per click. In a world where every second counts, this is incredibly significant.

“We had a big issue where the database became slower and slower as the times to click went up. Panellists’ times would go from 500ms to one, two or three seconds, and at that point they would stop completing questionnaires and just give up,” Carter explains. With spinning disks too slow to handle the transaction, the switch to flash has eliminated the issue. “It used to happen once or twice a week, but now we’ve removed spiking in performance like this completely, which is a really big thing from my team’s point of view.”

**MOVE TO PURE STORAGE SPEEDS-UP DATABASE TRANSACTIONS BY 4X**

Downloading data from completed surveys could be a challenge, explains Carter. “You can’t imagine how long downloading data from surveys with 100,000 completes would take, but now the download times are halved or more than halved,” he says. “End users don’t see it, but they have access to the data far faster.

“We’ve seen a definitive positive improvement. In the past, big projects were an issue – you might wait up to six hours to download data, with the analysis taking up the majority of that time. The new solution allows us to provide a more consistent service to the analysis staff, which previously could be very disruptive and cause user frustration. We are a lot more consistent and predictable now; I haven’t had a problem since we went live in January. Previously, there was a download performance issue every week.”



650 Castro Street, Suite #260, Mountain View, CA 94041  
 T: 800-379-7873 [www.purestorage.com](http://www.purestorage.com)