

Protecting the customer experience



[\[https://www.testingforschools.com\]](https://www.testingforschools.com) is the flagship digital assessment system operated by the education services provider GL Assessment (GL) [\[https://www.gl-assessment.co.uk\]](https://www.gl-assessment.co.uk).

GL launched Testwise as service in 2005 and it remains at the heart of their digital strategy. The service supports many schools in the UK and around the world, enabling them to assess and determine their pupil's educational needs. In the UK, September is an ultra-busy time for Testwise, when schools assess pupils at the start of the academic year.

The importance of quality of service

It is critical to GL's reputation that Testwise is available and error-free, especially during key periods, like September. Any downtime or failure in these periods is extremely disruptive to the schools using the system. Disruption that leads to dissatisfied customers, bad publicity and reduced uptake of the digital option.

Testwise 2010

In preparation for the September 2010 period, Testwise had been subject to a re-platforming and re-engineering exercise designed to address issues of performance, reliability and scalability. SQC supported this exercise, throughout 2010, delivering extensive load and failure testing that drove the engineering team's priorities. The results were significant improvements in the available quality of service and a highly successful September 2010. See the separate case study Testwise 2010 [\[https://www.sqc.co.uk/assets/documents/sqc-case-study-testwise-2010-2.pdf\]](https://www.sqc.co.uk/assets/documents/sqc-case-study-testwise-2010-2.pdf) for the full story.

Protecting their investment

The board of GL recognised they had gained significantly from the re-engineering of Testwise. They could now pursue a 'Digital First' strategy, confident that their platform would, not only, provide a great and dependable customer experience, but that it would support their ambitious plans for expansion of digital uptake both within and outside the UK.

The board also recognised that a key factor in the success of the exercise had been the involvement of SQC as a provider of specialist technical testing services. They understood that this was not a capability they could ever develop and sustain in-house. A first-rate in-house development team, yes, they had that, but they knew that first rate technical testing is the remit of specialists.

The board also knew they could not take their eyes off the ball. They knew technical testing, especially load testing, needed to remain a part of their normal engineering processes. Not something done on special occasions, but something done as a matter of course before making any significant change. Naturally, given their experience on Testwise 2010, they turned to SQC to shape and provide this service.

A standing load test capability

There are numerous approaches that could have been tried to meet GL's need. The one settled on was a fully managed, long term, load testing service provided by SQC. This has been proven to be a highly effective approach that has served both GL and SQC well since its inception at the end of 2010.

GL have an ongoing requirement for flexible load testing

GL wanted to ensure load testing became a standard part of their development culture. Something that was recognised as a core practice and not a 'nice to have'. They also recognised that this was not a competency that they possessed or could establish in-house, at least not to the same standard that they had become used to during the Testwise 2010 load testing. The board recognised that working with an external supplier, SQC, was clearly the best solution for their needs. Yet, despite not being in-house, the service needed to be 'on tap', able to respond rapidly to changes in plans and developing situations. It also had to be effective, a low burden on the development team, efficient and value for money.

Best met by a managed service

It was clear to both GL and SQC that the best solution was going to be a long term managed service load testing capability. One where, from a GL perspective, things happened 'by magic'. They did not want to be concerned with how the load testing implemented, they did not want to be concerned with the management and operation of load generator infrastructure. They did not want the challenges of finding, recruiting and retaining engineers with the skill to deliver load testing. What they wanted was an outcome on demand.

The managed service approach was best both parties. Concerns that might arise in some cases over such a relationship were not a barrier because of the high, mutual, levels of trust that exist between GL and SQC. Trust in all spheres including, technical, commitment and commercial. The continuity and length of the arrangement allowed SQC to invest time in optimising the delivery, for both efficiency and responsiveness, reducing the whole life cost for GL and providing the seamless 'on demand' integration with their development activities they required.

Service History

The service arrangement came into existence in the final quarter of 2010, following the Testwise 2010 campaign and the 2010 September busy period. GL have renewed this arrangement twice; the history of this arrangement is shown below.

First Service Period	A two-year plus, business as usual managed service covering Q3/2010 through to Q4/2012
Second Service Period	A three-year managed service including business as usual operations and replacement of the legacy load test tool with a new testing framework able to support testing at increased levels of load. Covering 2013 through to the end of 2016
Third Service Period	A further three-year service arrangement covering business as usual operations for the recently launched Azure based replacement of the Testwise technology stack

Value for Money

The service provided GL with on demand testing from a highly capable team for less than the cost of employing one full time engineer. Excellent value for money, especially as this included all tooling and operational costs.

Seven years successful of service

The service has been operating for seven years and looking back over that period there are some key, experience based, points that stand out.

Service capacity alignment with the capacity of the engineering team

It is hard to know what will happen in two or three next month, let alone what will be happening two or three years down the line. When the service was established it was based on a reasonable view of what would happen between then and September 2011, but that was it. There was no real view beyond that point.

One thing that was known was the capacity of the development organisation. There was a reasonable measure of how much 'output' they could achieve per year. The approach taken was to align the load testing service capability and capacity so that it was in balance with what the development team could produce. This proved a sound judgement for, although there have been some significant transformational changes within GL and to Testwise during the seven-year period, at no point has the service fallen short of requirements, nor has the service 'envelope' been breached significantly enough to warrant a change request.

Service flexibility was very important

The service was conceived as a capability to support the evolutionary development of the Testwise platform, the one tested during 2010. The reality has been more radical changes than were envisaged when the service was conceived. Despite this, the arrangement has been flexible enough and has been refined fast enough to absorb these events.

The major 'surprise' came with the purchase by GL of a separate organisation, The Test Factory [<http://www.thetestfactory.com/>], an organisation that was developing an Azure based platform for delivering digital testing services.

This purchase facilitated a programme to move the Testwise service from an on-premises J2EE based application to an Azure hosted Dot Net application. Needless to say, given the exceptional September load profile that hits Testwise together with the heavy historic investment in the non-functional engineering of the J2EE java application, this was not a trivial exercise. The full extent of the challenge really emerged during the first load tests of the new platform.

It took fifteen months of hard work, non-functional testing and engineering, to get the replacement system to the state where migration from the J2EE application occurred. Fifteen months of testing that SQC delivered within the established service envelope.

Over the seven years major activities of the service have included:

- Testing ahead of every September period. The load tested at has risen by a factor of nearly six over the load tested in 2010. SQC has enhanced the load generation capacity year on year to support this.
- Moving all load generation to the Amazon AWS service, replacing the on premises load generators provided and operated by SQC.
- Testing of the introduction of replication into the on premises Oracle database tier as well as a full-scale database upgrade.
- Replacement of GL's obsolete load testing tools with a custom, cloud based, load simulation solution.
- Creation of the load test capability for the Azure Testwise solution. Replicating the work done to develop testing of the original Testwise solution for the new solution.
- Extensive testing of the Azure solution. Campaigns that ran over fifteen months and that involved nearly four hundred extended load test runs during that period delivering millions of tests and other activities on to the system. Whilst, in parallel, supporting the production Testwise system.

Why a continuous service?

Both GL and SQC know that a continuous service, the Testwise Service Protection model, is the best option to ensure load testing is an established, effective, part of the development process. Why is this so? Why would an organisation choose to go into a long term managed service relationship when they could contract work on a case by case basis, seeking to optimise each piece of work separately?

Benefits

A long term managed service solution had major benefits over a case by case approach including:

Avoiding Commercial Delays

Attempts to procure service as a series of separate projects, each with its own commercial arrangements, slow things down and create overheads. Each case requires analysis, estimating, a proposal and approval, increasing costs as well as adding delays.

No Loss of Readiness / Start-up Delay

Being ready to test at short notice is important. With a project-based arrangement readiness is inevitably lost between periods of activity. When the time comes to start again there is extra work recalling what is required and a delay as work now must be done to align the load test tools and scripts with the latest version of the system.

A continuous, managed service, approach puts the onus on the supplier to maintain readiness and avoid delays. It also creates the opportunity for the supplier to remain ready, something they do not have if they are 'off the project' with no access to systems.

Efficiency

Avoidance of the overhead of repeated commercial activities and of the costs incurred when 'starting from cold' means that money spent is spent on value adding activities. There are more 'bangs for their bucks' making activities more cost effective.

Responsiveness

A development process is likely to create many different types of demand for load testing. Some predictable, test the next major release, others not so, 'we have had a problem reported by a customer can we test it?'. A standing service provides flexibility and an ability to respond quickly to demands.

Comprehensiveness

The model provides the ability to execute very small but important tasks. Ones that would not be done if there was a need to make special arrangements to get the task done. This helps to ensure that testing is part of the fabric of what is done, not something for big, ceremonial occasions.

Guaranteed Supply

Inevitably, on any system a supplier who has been involved in its development has a lot of specialist knowledge and a capability that an alternative supplier would take a long time to build. By entering into a long-term agreement, clients ensure the continuing availability of these assets, whenever they need them. They eliminate the risk that they will need to move to a new supplier and go through, and pay for, the development of similar knowledge and capabilities.

How the Testwise service works

Two key objectives have driven the way the managed service was established and has evolved:

1. Avoiding any burden on GL to do with the maintenance and operation of the services.
2. Allowing SOC the flexibility to provide the service in the way they deem appropriate.

Partnership

GL trust SOC. Trust SOC technically, trust SOC's motivation and trust SOC commercially. The two organisations work together in a true long-term partnership arrangement.

Well-structured remits

The remits of the two organisations are clearly understood. The allocation works well and has allowed load testing to operate in a highly efficient and effective way.

GL's remit

- Provision and operation of the Testwise test environment(s) used for load testing
- Importing of SOC generated test data sets and executing test data roll-backs
- Investigation of functional anomalies detected and reported by SOC.
- Assistance in clarifying browser to server API call operations and dependencies
- Platform and application-internal operational data collection during tests
- Investigation and remediation of performance and reliability issues

SOC's remit

- Conceptual ownership and implementation of the load generation tool chain
 - Conceptual ownership and implementation of the test data tool chain and pipeline
 - Provision and operation of the load generation platform used for load testing
 - Generation of test data sets for loading via data import into the test environment
 - Tactical, via the end customer interface, generation of supplementary test data
 - Work load analysis and synthetic load definition
 - Maintenance of load simulation compatibility with the Testwise implementation
 - Alignment of simulated user and browser behaviour with the 'real-world' behaviours
 - Selection of test characteristics to suit the risk to be addressed
 - Preparation and execution of tests
 - Collation and analysis of system behaviour under tests
 - First line investigation of anomalies to eliminate false positives
 - Overall non-functional health assessment
-

A regular cadence of test cycles

Test campaigns are regular, with a maximum gap between campaigns being enforced. Great importance is placed on maintaining this cadence. The reason? Work that is repeated regularly becomes predictable and more efficient. Work that is performed sporadically with long gaps between cycles becomes error prone, unpredictable and less efficient. The more complex the work the more this is the case.

Load testing is complex, there are so many things that must be right for it to work. Many times, people get something wrong. In this situation it can take many attempts to achieve a result, something a practiced team would have done in one cycle. Failed test attempts lead to incorrect results and / or increased testing costs. A regular pattern of testing saves money and improves the quality of the service.

Close collaboration

Most of the decision making and collaboration occurs at a day to day engineering level. What is to be tested and when it is to be done is dealt with at this level. Practical coordination occurs at this level, problem solving occurs at this level. There is little management overhead, to all intent and purposes the service is 'part of the fabric' and 'stuff just gets done'.

The service is a 'partnership of equals', both parties have an equal say in what needs to be done and in what order they should be done. Demands for testing do not only come from the GL team. The SQC team may need to run tests to validate changes to the load test framework. This need creates a demand on GL, a demand that they stand up the load testing environment to enable this and that they support with investigations.

A valued partnership

SQC and GL have had a long partnership. Though focussed on the non-functional testing of Testwise, in tow incarnations, GL understand the broad capabilities of SQC and have called on these when necessary. Additional activities SQC have performed for GL include:

Software Development Review	At the start of 2011, prior to GL's decision to purchase a separate development company, SQC brought its experience of software delivery management into play to perform a strategic review of GL's internal software development organisations
System Options Assessment	SQC undertook an analysis for GL of the options available to consolidate the production of analytical reports for both paper based and digital assessment onto a single reporting platform
Automated Testing of Report Equivalence	Testwise's rule-based answer scoring and report generation is complex. The move to Azure was high-risk. Combinations of tests, answer values, reports and report options were astronomical. 'Smart' testing was needed. SQC chose to replay production answer sets through both reporting chains and to compare the outputs. This solution included using Selenium to implement simulated pupils to feed answers via the browser where preliminary scoring occurred.
Testing of a new Digital product	A digital version of an old desktop product needed testing to ensure it provided the same analysis as the old one, an Excel workbook full of formulas and macros with little documentation. Testing the inhouse test team were too busy to deliver and so GL turned, without hesitation, to their partner to get this work delivered.

GL's experiences of and opinion on SQC is neatly summarised in the statement...

"Would I recommend SQC to other customers? Without a shadow of a doubt, probably the most recommended and trusted partner I have the pleasure to work with in over twenty years in the industry."

**Paul Webster, GL's
Operations Director
2010**

Need to know more?

To find out more about technical testing services, including load and performance testing, provided by SQC please visit the technical assurance section of our website [<https://www.sqc.co.uk/perform/assurance/technical-assurance>] or email our enquiry mailbox [enquiry@sqc.co.uk] and we will get back to you.