LOADBALANCER

WHITE PAPER

The practicalities of delivering ‘cloud first’ in the NHS

Discover how to use load balancing to reduce commercial and technical risks when migrating vital healthcare IT systems to the public cloud

Executive Summary

‘Cloud first’ is a core policy of the UK Government. Public sector organizations, including NHS hospital trusts and other health and social care organizations, are therefore being urged to proactively consider migrating their data, IT services and applications to public cloud platforms wherever it is feasible to do so.

Why cloud first in the NHS?

The rationale behind this policy is that the use of public cloud platforms will help the NHS to deliver its ambitious digital transformation strategy, by capitalizing on these potential cloud benefits:

1. Scalability and flexibility to adapt and grow their IT services.
2. Reduced strain on their in-house IT resources.
3. Free up space currently used for on-premise data centers.
4. Improve their budgeting by moving to an OPEX model of funding.
5. Improve application interoperability, to deliver more patient-centric services.

What about data security?

The practicalities of migrating to a public cloud platform are complex. Understandably, many healthcare organizations are reluctant to migrate their data, IT services and applications to the public cloud for the following reasons:

1. Concerns about cybersecurity attacks and data breaches.
2. Concerns about data governance.
3. Concerns about data regulations and compliance.

So much so that the government has had to issue clear guidance to the sector, and offer risk assessment frameworks to help healthcare organizations mitigate security risks in cloud migrations.

What are some of the commercial and technical risks of a cloud first approach?

1. ‘Bill shock’ due to unpredictable and increasing costs.
2. Poor application performance in the cloud.
3. A shortage of specialist skills in existing IT teams.
4. Cloud provider ‘lock-in’, making it hard for them to switch in the future, if costs get too high or priorities change.
5. The unavailability of critical applications in the event of an outage.
6. Many applications may be too complex to simply ‘lift and shift’, requiring replatforming or rebuilding.
7. Data corruption and performance degradation during the migration period itself.

> *“NHS organizations that are currently thinking ‘cloud first’ should, therefore, think load balancers second.”*

What role do load balancers play in the cloud?

Load balancers play an essential role in cloud platforms, managing the flow of data to the cloud and directing all user traffic to the server that will deliver the fastest response.

NHS organizations that are currently thinking ‘cloud first’ should, therefore, think load balancers second.

How do platform-agnostic load balancers mitigate risks?

While public cloud providers offer their own load balancers as part of their own cloud solution, these native load balancers can, in fact, increase the technical risks in hybrid environments because they do not ‘talk’ with other load balancers on other platforms, without complex and costly integrations.

By contrast, having a single, platform-agnostic load balancing platform that works across all local and cloud services, can help unify these components and mitigate many of the commercial and technical risks - particularly those in hybrid environments.

What should you look for in a platform-agnostic load balancer vendor?

Doing thorough due diligence on platform-agnostic load balancing vendors could make all the difference to the success of your migration to the public cloud, and the long-term delivery of your digital transformation goals.

It is prudent to consider:

1. How a platform-agnostic load balancing platform can be integrated for use across on-premise, legacy data centers and public cloud platforms, and to mitigate your many commercial and technical risks.
2. How easy the load balancer is to configure, use, manage and update.
3. How to avoid vendor ‘lock-in’ and take advantage of flexible licensing options that work for your budget.

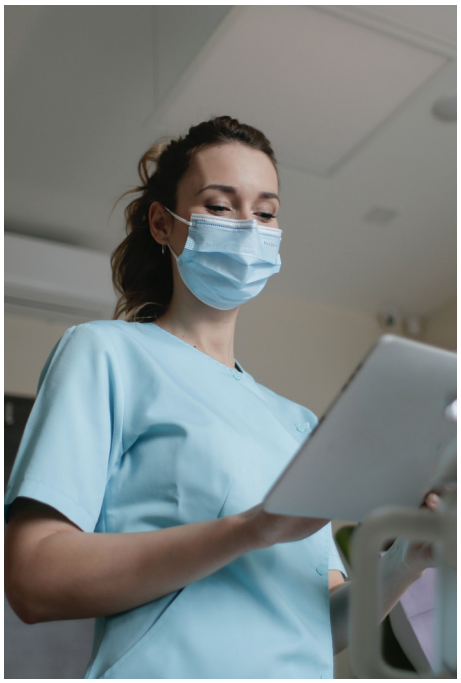

Why cloud first in the NHS?

The UK government introduced 'cloud first' in 2013, directing that all public sector bodies "should consider and fully evaluate potential cloud solutions first" before adopting other options[1]. It is one of six key architectural principles outlined in the UK Government's policy paper 'The Future of Healthcare – Our Vision for Digital, Data and Technology in Health and Care'[2]. It makes the assumption that all NHS healthcare IT services should eventually move to the cloud, replacing all locally managed services.

This is a somewhat utopian view however, at least in the short term. In practice, health and social care providers are unlikely to migrate all their IT services and data to the public cloud for a variety of reasons. Nonetheless, even a partial migration is likely to deliver many of the following benefits:

Improved flexibility and scalability

A public cloud platform can give healthcare organizations greater flexibility to grow existing IT services, and develop new ones, on demand, in response to emerging and unexpected situations. This supports the innovation and deployment of new technologies such as artificial intelligence and high performance computing.

Reduced strain on in-house IT staff

Keeping life-saving medical applications secure and up-and-running 24/7 places a lot of responsibility on in-house IT teams. Having a public cloud partner shares this responsibility, while also improving resilience and back-ups.

More space to deliver patient services

Anyone who has seen patients on trolleys in corridors and store-rooms emptied to make room for beds knows that space is at a premium on hospital sites. By using the cloud, NHS trusts can free up on-site data center facilities for other patient-centered services.

OPEX not CAPEX expenditure

A migration to the public cloud is, in effect, a move away from capital expenditure (CAPEX) to operational expenditure (OPEX). Rather than having to specify, procure, buy, build, manage and maintain their own IT equipment, organizations consume the exact services they need.

Secure, shared access to patient data

When data and IT services are all available centrally, via the cloud, healthcare professionals can access them remotely, with appropriate security access controls. Consequently, it becomes easier for them to deliver more patient-centric services within communities and at patients' homes.

1 Government Cloud First Policy [Internet]. GOV.UK. [cited 2021 Nov 29]. Available from: https://www.gov.uk/guidance/government-cloud-first-policy

2 The Future of Healthcare: Our Vision for Digital, Data and Technology in Health and Care [Internet]. GOV.UK. [cited 2021 Nov 29]. Available from: https://www.gov.uk/government/publications/the-future-of-healthcare-our-vision-for-digital-data-and-technology-in-health-and-care/the-future-of-healthcare-our-vision-for-digital-data-and-technology-in-health-and-care

What about data security?

Data security is the primary issue that concerns all organizations when they are considering a migration to the public cloud, regardless of industry. For decision makers in the health and social care sector, however, this concern is particularly prominent.

Hospital trusts and healthcare providers generate, handle and store vast amounts of highly sensitive personal data about patients, and any data security breach could potentially be very damaging. Healthcare organizations therefore need to ensure that rigorous security measures are put in place and enforced to protect their data – regardless of whether that data is held on-premise, or in the cloud.

> *"Rigorous security measures are needed - regardless of where the data is held."*

While healthcare providers may be nervous about the data security implications of using the public cloud, the latest guidance, issued jointly by NHS Digital, the Department of Health and Social Care, NHS England and NHS Improvement, clearly states that 'NHS and social care organizations can safely locate health and care data, including confidential patient information, in the public cloud.'[3]

There are, however, some conditions. In brief, the advice from NHS Digital can be summarized as: understand the data you are handling; assess the risks associated with this data; and implement proportionate controls.[4] Healthcare organizations can use standard risk assessment frameworks, such as NHS Digital's Health and Social Care Cloud Risk Framework to gauge the security requirements for different types of data and then make informed decisions about whether public cloud platforms provide sufficient protection on an application by application basis.

The very fact that healthcare organizations are most concerned about data security means that the security implications of a cloud migration will be closely examined and risks mitigated. Senior Information Risk Owners (SIROs) within healthcare organizations will ensure that they are satisfied with the security arrangements offered by public cloud providers, implement strong data governance policies and ensure compliance with regulatory conditions,

3 NHS and Social Care Data: Off-Shoring and the Use of Public Cloud Services [Internet], NHS Digital. [cited 2021 Nov 29]. Available from: https://digital.nhs.uk/data-and-information/looking-after-information/data-security-and-information-governance/nhs-and-social-care-data-off-shoring-and-the-use-of-public-cloud-services

4 Health and Social Care Cloud Risk Framework [Internet], NHS Digital. [cited 2021 Dec 1]. Available from https://digital.nhs.uk/data-and-information/looking-after-information/data-security-and-information-governance/nhs-and-social-care-data-off-shoring-and-the-use-of-public-cloud-services/cloud-risk-framework

such as the General Data Protection Regulation. There are, however, many other risks associated with the use of public cloud platforms that do not relate to data security, and these can easily be overlooked, with huge commercial and technical consequences.

What are some of the commercial and technical risks of a cloud first approach?

1. Bill shock

It can be incredibly difficult to predict the real costs of cloud computing, as organizations need to estimate not only what their data volumes will be, but also how they will need to access data and the level of processing power that will be required. As a result, some public sector bodies in the UK that rushed to migrate all their systems to the public cloud experienced what is commonly referred to as 'bill shock'. They are now repatriating some of these systems back to on-premise platforms to try to better manage and reduce their costs.

2. Slow application performance in the cloud

In the healthcare sector, file sizes can be very large and transferring them up and down, to and from the cloud, could lead to slow application performance. The file size itself is dependent on a number of variables, such as whether it is a 2D or 3D image, the dimensions, and the resolution. As a rough guide, a digital mammogram can be 50MB and a digital radiography image can be 32MB.5 When it comes to CT and MRI scans, the individual scans can be quite small, but a study or exam consists of 100s or even 1000s of scans (slices) and the DICOM file contains all of these together. So the size of the DICOM file for a CT set ranges from 300 MB for a plain brain study, to 1GB for a contrast study of the abdomen, while an MRI scan can be up to 250 MB.6

For radiologists currently battling to catch up on the huge backlog of work as a result of the pandemic, any delays in receiving images will be a significant obstacle in their day-to-day work. For this reason, migrating imaging systems to the cloud may not make sense if performance is in any way impaired.

5 J. A. Seibert, PhD, FSIIM; University of California, Society for Imaging Informatics in Medicine. Available here: https://siim.org/page/archiving_chapter2
6 A. Chandrasekhar, MD , MBBS Radiodiagnosis, Rajarajeswari Medical College and Hospital. Available here: https://www.quora.com/What-is-the-file-size-of-a-3D-image-generated-using-CT-Scan

3. Skills shortage

As the National Audit Office cautioned in its 2019 guidance for audit committees on cloud services, "the cost and effort of moving to cloud computing solutions, and the skills required to manage them effectively should not be underestimated."[7] The challenge will be magnified if healthcare organizations use a combination of public cloud platforms and on-premise IT infrastructure – and don't take steps to adopt a standardized approach to load balancing.

4. Platform 'lock-in'

"There are risks with a cloud first approach that need to be mitigated"

If healthcare organizations move all, or most, of their IT services to one public cloud platform, this makes it much harder to change provider in the future if service levels are not met, costs rise, or they want to pursue a multi-cloud strategy. Lock-in is one of the risk factors highlighted in NHS Digital's Health and Social Care Cloud Risk Framework, published in October 2021. As the report explains, flexibility may be adversely impacted by the adoption of a specific public cloud provider's unique services.[8] As flexibility is one of the desired benefits of pursuing a cloud strategy, this potential pitfall should be avoided.

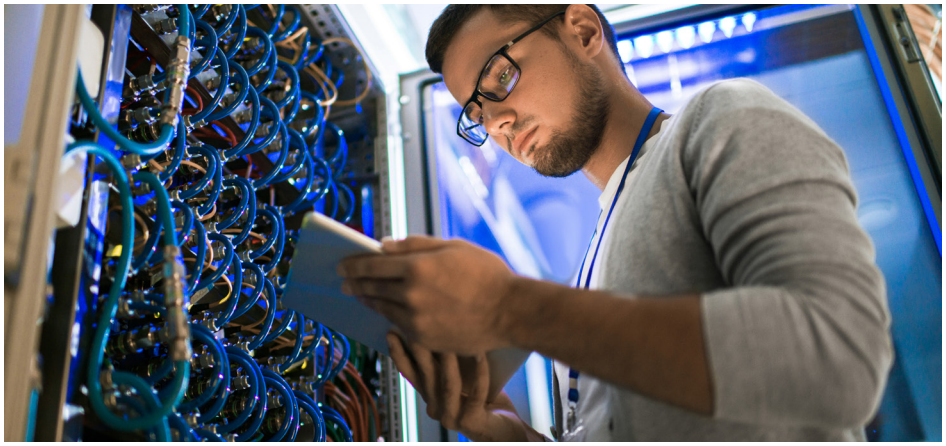

5. Unavailability of critical applications

If a clinician cannot access vital patient information in an emergency situation, a patient's life could be put at risk, and while cloud outages are rare among the major providers, they do still happen. Internet connectivity to a cloud can also fail. It is hardly surprising then that another of the risks highlighted in NHS Digital's Health and Social Care Cloud Risk Framework is the risk that critical applications could become unavailable. As the report explains, 'Network connectivity to the cloud becomes a critical dependency and there is a risk of introducing a Single Point Of Failure (SPOF).'[9]

7 Guidance for Audit Committees on Cloud Services, National Audit Office, April 2019. Available from: https://www.nao.org.uk/wp-content/uploads/2019/04/Guidance-for-audit-committees-on-cloud-services.pdf
8 Health and Social Care Cloud Risk Framework [Internet], NHS Digital. [cited 2021 Nov 29]. Available from: https://digital.nhs.uk/data-and-information/looking-after-information/data-security-and-information-governance/nhs-and-social-care-data-off-shoring-and-the-use-of-public-cloud-services/cloud-risk-framework
9 Health and Social Care Cloud Risk Framework [Internet], NHS Digital. [cited 2021 Nov 29]. Available from: https://digital.nhs.uk/data-and-information/looking-after-information/data-security-and-information-governance/nhs-and-social-care-data-off-shoring-and-the-use-of-public-cloud-services/cloud-risk-framework

6. Too complex to 'lift and shift'

Moving applications to the cloud isn't always as straight-forward as 'lift and shift.' Many applications will require a degree of re-platforming (re-architecting the application to work in a cloud environment) or even rebuilding as cloud native applications to enable them to work in the cloud. In its report 'Digital Transformation in the NHS', the National Audit Office reported that moving NHS Spine to the cloud 'would be unlikely to save money and would involve a significant effort'. As a consequence, this key health and social care IT infrastructure will not be moved to the cloud for the time-being and instead will need to be completely redeveloped to transform it into a feasible cloud-based system.[10]

7. Disruption to services during the migration

The migration to the cloud will inevitably create risks, as even with extensive testing, organizations can never be entirely sure that applications will perform seamlessly in a cloud environment once they go live. Firstly, there is a very real risk of data corruption or data loss and, secondly, there is the risk that the data migration will take too long, with interoperability issues adding to these delays. Organizations will therefore need to plan their migrations carefully and ensure that they can push traffic back to the original on-premise environment, very quickly, if any degradation in performance is observed.

10 Digital transformation in the NHS, National Audit Office, May 2020. Available from: https://www.nao.org.uk/report/the-use-of-digital-technology-in-the-nhs/

What role do load balancers play in the cloud?

Load balancers are essential in cloud environments. They distribute user traffic across multiple servers in the cloud, to prevent single servers from becoming overloaded and ensure that users all experience the same high performance. When demand for an IT service grows or there is a sudden unexpected increase in traffic – such as to a healthcare website – it is the load balancer that will direct traffic to additional servers to maintain service availability.

Native load balancers

Public cloud providers offer their own native load balancers as part of their cloud solutions, typically offering different types of load balancers for different functions. AWS, for example, provides a choice of application load balancers, network load balancers, classic load balancers and gateway load balancers. Each of these proprietary load balancers is designed to handle different types of traffic and support different network protocols. Consequently, healthcare organizations that opt for the native load balancing approach will invariably need multiple, different load balancer instances within a single cloud platform to provide their required application functionality.

As we have discussed here previously, it is not feasible for most healthcare organizations to move all their applications and IT services to the cloud. It may also not be prudent for healthcare organizations to move to a single cloud platform for fear of lock-in or bill shock. Consequently, the trend in the sector currently is to use one or more public cloud platforms, alongside on-premise IT equipment. This requires sound decisions to be made about what data needs to be kept locally, what data will perform better in the cloud, how the migration should take place, and in what timeframe.

In these hybrid environments, load balancers play an even more critical role.

> *Without a platform-agnostic load balancing platform, flexibility is therefore significantly constrained.*

Platform-agnostic load balancers

Organizations can create a platform-agnostic load balancing platform, using load balancing technology from a single vendor that will interface with on-premise servers, and servers across multi-cloud platforms. Using global server load balancing (GSLB), a feature available in many load balancers, organizations have the flexibility to direct traffic to the most appropriate location; for example, an on-premises server, a remote data center, or resources in the cloud.

Without a platform-agnostic load balancing platform, flexibility is therefore significantly constrained.

Native cloud load balancers from Microsoft, Amazon, Google and other smaller cloud providers are incompatible both with each other and with on-premise and legacy load balancers. Consequently, if an organization has two different sets of load balancers for two cloud platforms and another load balancer at its data center, it will not be able to easily distribute workloads across all three platforms without costly work-arounds and integrations.

How do platform-agnostic load balancers mitigate these risks?

Platform-agnostic load balancers balance traffic in a standardized way across multiple cloud platforms and on-premise infrastructure, mitigating the commercial and technical risks previously discussed.

> *Public cloud outages won't automatically result in the redistribution of traffic without intelligent load balancing.*

1. Managed, phased migrations

When the same load balancing technology is in place in the on-premise environment and cloud, they have more control over how they manage their migrations to the public cloud, reducing the risk of unanticipated performance issues. The IT team can make use of GSLB to send a controlled amount of live traffic (such as 10%) to AWS, for example, and then monitor how well things are working. Assuming there are no problems, the IT team can then gradually increase traffic over time, in a managed way, until such time as 100% of traffic is going to AWS. If any issues are experienced, the load balancers can instantly reroute traffic back to the on-premise data center, while glitches are ironed out.

2. High application availability

Although reasonably rare, outages in public cloud platforms do occur. In such cases, organizations that have a common load balancing platform can instantly and imperceptibly reroute all their traffic from the unavailable cloud environment to an alternative cloud provider or on-premise infrastructure. This redistribution of traffic will not occur automatically unless the same load balancing technology is used across all cloud and local services and can make intelligent decisions about where to direct traffic to ensure users experience consistent high performance.

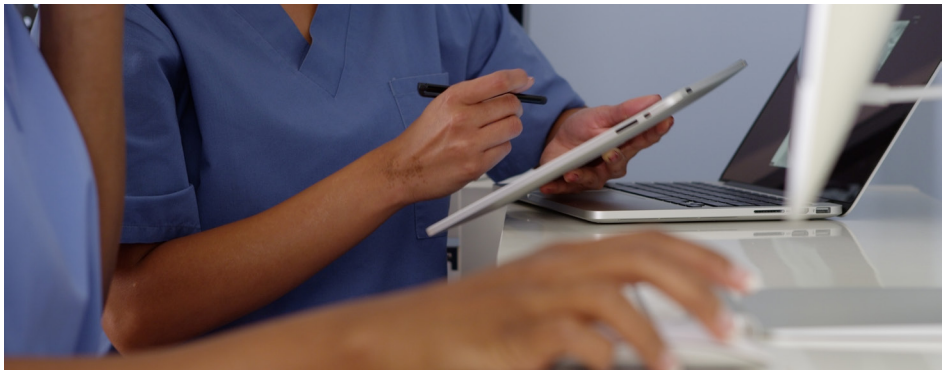

3. Improved cost control

Native load balancers, offered by cloud providers, are typically charged based not just on data flow but also on how data is processed. This makes them more expensive in the long run. Given how complex it can be to accurately project the cost of running services in the cloud, many healthcare organizations find it reassuring to have control over what proportion of traffic is serviced via the public cloud by using load balancers. If one particular cloud platform gets too expensive, they can redirect traffic to an alternative, cheaper cloud provider or move workloads back to on-premise data centers.

4. Flexibility for the future

Rather than being 'locked in' to one particular cloud platform, a common load balancing platform gives organizations the flexibility to migrate services and applications to alternative cloud providers on demand (and switch back again easily if necessary). The leading cloud providers have differing strengths. AWS, for example, is renowned as a strong platform for DevOps projects. The reasons why an organization might select one cloud platform over another may change in the future, perhaps due to changes in government policy or data regulations. A separate load balancer will keep its options open.

5. Easier IT management

Having a single load balancer technology that the IT team can use for on-premise and cloud deployments will significantly simplify IT management. An organization with its own data center and two cloud providers would not have to train its IT team in the use of three different load balancing technologies. It would also be easier to recruit new staff and get them up-to-speed as there would be a common load balancing framework across local and cloud-based systems.

What should you look for in a load balancer vendor?

Cloud–agnostic solutions

An organization may already have decided upon AWS, but what if, a few years down the line, it needs to change. Using a cloud-agnostic load balancer will enable it to switch between Azure, AWS, GCP and other cloud platforms on demand, without having to invest in new load balancing technology or retrain staff.

License flexibility

Organizations that have on-premise load balancers now and are planning to move to the public cloud should find out if they can use their existing solutions. For example, Loadbalancer.org offers a 'Freedom License' allowing customers to move their existing license whether hardware or virtual to AWS, Azure or GCP cloud platform.

Experience in healthcare

Healthcare is a sector like no other. Data is incredibly sensitive and application performance is paramount. Health and social care organizations will need a load balancing vendor that understands this and has long experience of load balancing critical patient care systems. Loadbalancer.org has, for many years, been the preferred healthcare solutions partner of several key vendors of healthcare solutions, including Philips and Fujifilm.

Simple management

Already under pressure, IT teams in hospitals and other healthcare settings need IT solutions that are easy to deploy, manage and maintain. Having a single load balancer for use across on-premise and cloud platforms will significantly reduce complexity, particularly in multi-cloud environments. It is advisable to look for vendors that have load balancers for AWS, Azure and GCP – plus a commitment to customizing solutions for emerging new cloud platforms.

Superfluous functionality

Many load balancers on the market today are packed full of features, making them both costly and complicated to implement. When NHS funding is tightly constrained, organizations do not need to spend money on high-end products with superfluous functionality that they will never use. GSLB is an essential function in public cloud environments; products with other, more advanced features should only be considered if they will deliver long-term value for money.

Consultative approach

Moving to the public cloud is a big step and NHS organizations are likely to need support along the way. The large public cloud providers all offer support with load balancing, but only for an additional cost, and these ‘hidden fees’ are not included in the online cost calculations that are found on their websites. For peace of mind, it is worth looking for a vendor that provides a consultative approach to support and works with its customers to ensure the success of their cloud strategy, at no additional cost.

About Loadbalancer.org

Loadbalancer.org's mission is to ensure that its clients' businesses are never interrupted. The load balancer experts ask the right questions to get to the heart of what matters, bringing a depth of understanding to each deployment. Experience enables Loadbalancer.org engineers to design less complex, unbreakable solutions - and to provide exceptional personalized support.

LOADBALANCER

Visit us: www.loadbalancer.org

Phone us: +44 (0)330 380 1064

Phone us: +1 833 274 2566

Email us: info@loadbalancer.org

Follow us: @loadbalancer.org