



How load balancers solve key problems and accelerate application delivery

Healthcare IT



COVID-19 has taught the IT industry a valuable lesson: worst case scenarios really do happen.

“ *When the coronavirus emerged in early 2020, it not only seriously disrupted healthcare services; it also brought about a seismic shift in application usage. Clinicians began working from home, dependent on collaboration apps and remote access services to continue to fulfil their roles. At the same time, patient demand for telehealth services increased dramatically. Application usage soared, as did patient expectations for 24/7 healthcare services - yet many IT teams had little or no access to their data centers.*

The healthcare organizations that coped best in these extraordinary circumstances were those that had prepared for a worst case scenario, expecting it never to happen. They had load balancers in place, capable of managing the unprecedented surge in application traffic and maintaining uninterrupted services. Throughout national and local lockdowns, load balancers worked behind the scenes to keep vital collaboration tools, critical applications, platforms and communications services available for an unprecedented number of users.

As we learn to live in new uncertain times, we cannot let our guard down. Disasters of all kinds can happen at any time. Next time, it might not be a disease, but a fire at a data center, a terrorist alert, or a cyber-attack. It's not what the disaster is that matters, but how ready you are.

Recent events have taught healthcare institutions that they need simpler load balancing solutions that they can set up quickly, manage remotely, and trust to deliver exceptional reliability for mission-critical applications to provide life-saving patient care. As everyone now knows, worst case scenarios can quickly become reality and we need to be ready for whatever happens next. ”

Malcolm Turnbull, Co-Founder, [Loadbalancer.org](https://loadbalancer.org)

EXECUTIVE SUMMARY

The healthcare IT ecosystem

Healthcare IT infrastructure is extensive and complicated. It needs to connect primary and secondary healthcare providers, hospitals, patients, drug companies, research institutions, pharmacists, IT teams, insurance providers and financial institutions.

This complex web of network, storage and computing resources are used, owned and maintained by hundreds, if not thousands, of different stakeholders - far more than could ever collectively be involved in a single project. So how can such an extensive IT ecosystem be modernized and future-proofed, when there are so many diverse needs and demands being made of it?

In this eBook we will answer this question by examining the following three solutions:

- Cloud migration
- Interoperability
- De-risking IT architecture

For each, we will explore the small but pivotal role load balancers play in delivering these solutions and mitigating some of the risks.





1. CLOUD MIGRATION SOLUTIONS

Leveraging platforms more effectively

Cloud infrastructure facilitates the collection, sharing, analysis and reporting of data from disparate software and data sources (e.g. telehealth, telemedicine, EHRs, EMRs, and billing systems). Ultimately this enhanced integration should make collecting and sharing data a quick and seamless experience for medical professionals - which results in better patient care.

The main appeal of cloud is that it can remove the local IT resource burden. Cloud infrastructure already exists and has the potential to remove significant strain on the hospitals own data systems, as well as enabling scalability and facilitating Artificial Intelligence (AI) and high performance computing. The caveat is that cloud migration is by no means a quick fix, and there are still workloads that healthcare institutions may wish to remain local for a variety of reasons.

Step One. Impact Analysis

Cloud migrations can be complex, with a lot of moving parts. Before defining your strategy, you will need to fully consider the following:

COMPLIANCE

It may be that data has to reside in the locality in which it is generated (e.g. on-premise, or in the country of origin). Nonetheless, some of the leading cloud providers align with HIPAA risk management programs and provide higher security standards that map to the HIPAA Security Rule.

SECURITY AND BUSINESS CONTINUITY

There is still the potential for data loss, even in the cloud (depending on the technologies, policies, controls, and services implemented to protect data, applications and infrastructure from threats). In healthcare it is critical that failover and failback are seamless when things go wrong, so services are not interrupted.

CALCULATING THE COST OF ON-PREMISE VERSUS CLOUD

It can be difficult to compare the cost of investing further in your own on-premise data center versus investing in the cloud (although tools do exist to help with this). The biggest challenge for data and planning is trying to predict what the data volumes will be, how you'll need to access it, and the level of processing needed.

USER IMPACT

Cloud migration should not have a negative effect on the end user. Patients and clinicians should ideally be able to access relevant data from anywhere, although there is clearly a balance to be struck between access and security.

SKILLS AND EXPERIENCE

Cloud is still relatively new, so there aren't enough experienced people currently in the market to meet demand, although lots of learning support is available.

Step Two. Agree a phased workload strategy

A clear short, medium and long term game plan and milestones are recommended. The focus here is on application stack performance and readiness.

In the short term....

REVIEW INDIVIDUAL APPLICATION READINESS

- What applications have you got?
- How are they architected?
- Do you need to lift and shift, re-architect or repurpose?
- How will you ensure application performance doesn't slow in the cloud?
- What workloads will you move initially e.g. long-term archives?

ASSESS RISK TOLERANCE

- What is your risk appetite?
- How will you make sure you don't lose data?
- How will you ensure you can maintain the benefits of the cloud?
- What's your cloud migration contingency plan?

In the medium term....

INCREASE WORKLOADS IN THE CLOUD

- Can you move more applications e.g. patient information and records?
- Some apps such as medical imaging will be much more challenging.

ADOPT DATA AND TECHNOLOGY STANDARDS

- How does this data need to be communicated from one place to another?
- Do you need to look at open standards of communication e.g FHIR?
- Where are the users?
- How do people store and access that information from anywhere?

In the long term....

In the long term you will probably look to move clinical applications to the cloud where possible, whilst maintaining a small local footprint. Exponential growth will need to be proactively managed to ensure storage, security, recovery and performance requirements remain fit for purpose.

Step Three. Determine your load balancing needs

For those new to load balancing, [jump to our website for a full explainer](#). For those already familiar with load balancing, read on....

Do you need to aid storage performance?

A complete cloud load balancing solution can optimize your storage environment. It can check the health of storage clusters, making decisions about where images or patient information should be stored. For example, is the data performing well? Can it get in and out of the cloud efficiently? How does it use multiple clusters of data?

Modern storage architecture is good at recovering from failure, rebuilding and ensuring that data integrity is preserved. But it doesn't account for maintaining consistent access for users, which is paramount in healthcare IT. Without a load balancer, if part of a storage cluster failed, yes the data would be preserved, but it might not allow users consistent access which could be catastrophic.

Do you need to reduce cloud costs?

Cloud can be a wolf in sheep's clothing. For example, although migrating to a large public cloud provider such as Amazon Web Services (AWS), Microsoft Azure or Google Cloud Platform (GCP) may be reasonably inexpensive, the costs quickly escalate once you're there.

A native load balancer within a cloud platform will calculate costs not just based on information flow, but also on how that information is flowing and how it is being processed. With a complete cloud load balancing solution however, running costs are much less because you're not paying for the processing of that data. You are only paying for the cost of the load balancer to run in the cloud.

Do you need intelligent decision making?

A complete cloud load balancing solution can also make decisions based on the content of a request. It can analyze the data and determine that it may not want to send that data to a particular place immediately, as that might be a less cost effective way of storing or using the information. So in this scenario, the load balancer helps build intelligence on the information flowing in and out of the cloud environment, and therefore the resulting cost implications.

Step Four. Configure your load balancer

The simplest way to explain how this might work is to provide an example of how Loadbalancer.org built a stable, scalable cloud-based platform for Fluid Networks to host a medical application, using our Enterprise AWS Cloud Solution.

[Refer to our website for this example of how to shift critical healthcare applications to the cloud to deliver scalability using a load balancer.](#)

“ *It's not just about the here and now - it's about where healthcare IT needs to get to.*

The future will bring sizable challenges in terms of architecture, data security and compliance, new technologies, and exponential data growth.

All of this will present complex interoperability problems, which will need to be overcome if IT teams are to simplify their networks and transition to a more sustainable working model. ”

James Loveday, Application Specialist, Loadbalancer.org



2. INTEROPERABILITY

Harnessing cross-functional workflows

Interoperability is one of the most important pillars of any healthcare IT ecosystem. Care needs to be coordinated across different settings, with access to electronic patient records and test results anytime, anywhere.

It is also a key pillar for AI and machine learning as it will facilitate data driven analysis to back up clinical decisions and speed up the time to diagnosis. The reduced administration achieved through interoperability therefore has tremendous potential to improve the patient and clinician experience.

But there are significant threats to achieving this. For example, layer-upon-layer of distinct teams exist across healthcare, using different types of data and resources in different ways, to do different things. So it's not just accessing the data that's a challenge - it's managing, analyzing and ensuring the high availability of the data using cross-functional workflows.

Step One. Determine your load balancing needs

Healthcare interoperability enables computer systems to connect, communicate, and exchange information with one another readily, even if they are on different platforms, used by different vendors. But that means the various software, networks, databases and applications need to exchange data and present that back to the user in a way that can be understood.

A large proportion of work that load balancers do is taking data from a modality (e.g. a CT or MRI image generated by a scanner), making a decision about where that image needs to be stored, and using a particular standard of communication with a health records system so it can be integrated. Because all this information traverses the load balancer, it is able to make intelligent decisions about where that data needs to go.

For instance, there are different methodologies that can be employed to support DICOM and non-DICOM workloads and their different applications - all of which a load balancer can help with. The load balancer helps mitigate the data risk of imaging information and metadata (and patient information and metadata) by managing these different workflows and the applications that sit on the back-end that require that information. The load balancer also manages the client side and modalities that sit on the front-end that are generating that information or drawing it back out (e.g. a clinician viewing images or patient records).

Step Two. Configure your load balancer

Again, you may wish to see an example of how Loadbalancer.org helped build a patient-centred smart healthcare ecosystem, ensuring interoperability.

[Refer to our website for an example of how to facilitate the interoperability of Fast Healthcare Interoperability Resources \(FHIR\) using a load balancer.](#)

3. DE-RISKING IT ARCHITECTURE

Simplify to reduce risk and costs

Simplifying IT architecture not only reduces costs and saves time, it also has the potential to improve the patient and clinician experience - providing efficient, remote access and ensuring data is not lost along the way, as well as protecting data security.

This is a challenge because historically healthcare systems have been very siloed, with individualized, unlinked systems. Hence there is a strong focus on building modular systems that support growth and can talk to one another. So the changing healthcare landscape demands that IT architecture be reviewed and reshaped to ensure it remains fit for purpose.

Step One. Architecture assessment

Before working out what needs changing, a thorough understanding of the existing patchwork of IT ecosystems is paramount:

- How are legacy technologies currently configured?
- How are new applications bolted on?
- Do the foundations need to be modernized first?
- Are existing resources being utilized effectively?
- Where are the workflow maps for each process and where do they cross?



Step Two. Determine your load balancing needs

Load balancing can help deliver agile digital architecture in a number of ways (many of which have been described on the previous pages). Here we explore workload migration and the facilitation of a per app approach.

MAP ALL SCENARIOS

For example, there may be an action to migrate a certain workload to the cloud, but there are some scenarios where that data might not be able to migrate. Therefore there needs to be a means of failing back, protecting you from odd events that may happen, causing potential data loss or outages. In this way, load balancers support a phased approach because they are able to help you move and test workloads, making sure they are fully operational prior to migration.

CONSIDER A PER APPLICATION APPROACH

Having a load balancer that is dedicated to a single application, only doing one job, has a number of significant benefits. The main one is that, with a dedicated approach, there is a smaller footprint. The load balancer is committed to that application and any changes are unique to it. This requires less resourcing for maintenance, and the load balancer remains responsible for that workload only. The result is that it is incredibly unlikely the load balancer can be overloaded by the addition of more services as needs grow. This de-risks the application and ensures the resources are always available to support it.

MULTI-SITE RESILIENCY USING GSLB

Using Global Server Load Balancing (GSLB) you can distribute traffic intelligently across server resources and sites i.e. physical data centres, hybrid, cloud, multi-cloud. So no matter where your applications reside, if there's any failure or part failure within that locality, your user is able to retain access, using the path of least resistance.

Just as a load balancer distributes traffic between connected servers in a single data center, Global SLB distributes traffic between connected servers in multiple locations (whether these servers are in an organization's own data centers or hosted in the public or private cloud) and are able to provide failover between these sites based on the availability and performance profile of the data centers. Should one server in any location fail, Global SLB re-routes traffic to another available server somewhere else in the world.

By enabling all user traffic to be switched instantly and seamlessly to an alternative data center in the event of an unexpected outage, Global SLB improves the resilience and availability of key applications, and enables organizations to constantly monitor application performance at geographically separate locations. It also ensures the best possible application availability across multiple sites. During routine maintenance, Global SLB enables organizations to temporarily direct user traffic to an alternative site in order to avoid disruptive downtime.

Note: please don't confuse Global SLB with geographic load balancing, which redistributes traffic dependent on internet-based geo-location.

Step Three. Choose a flexible load balancer

How you configure your load balancer will be unique to your software and application stack, existing degree of interoperability and IT architecture. One thing is a must though - your load balancing strategy will need to ensure future growth is supported adequately over the next 3-5 years.



4. WHY LOADBALANCER.ORG

Tech vendors and HCOs work with us because...

Healthcare IT specialists

Loadbalancer.org is the only load balancer vendor to have a division dedicated to the medical imaging and clinical workflow sector. For more than 20 years, we have worked with a range of 'blue chip' customers and partners such as Fujifilm, Philips, GE Healthcare, Change Healthcare, Carestream, Hologic & Agfa, and developed validated solutions for many others.

Our extensive experience engineering healthcare applications (from Enterprise Imaging to Clinical Workflows, EHR to Interoperability solutions) allows us to work closely with existing and prospective customers to develop solutions that facilitate high availability, scalability, and zero downtime for critical healthcare applications.

Part of the open source community

Our suite of appliances and virtualized load balancer solutions are all based on well-established open source technologies, including HAProxy, LVS and ModSecurity. We believe this makes them the most secure, flexible, efficient and powerful load balancers available. We are however not beholden to any one open source platform. This allows us the flexibility to cherry-pick the best components of each open source technology and also from different technologies. All of this while providing the support services, documentation and SLAs that you would expect of a proprietary solution.

A unique and refreshing approach

We pride ourselves on developing fully featured, high performing products, that are easy to configure, use and manage - backed up by a top-notch, tireless support desk. At Loadbalancer.org, we want you to be in control, so we don't do end-of-life. So we will always support your product - and that is a promise. Our products run on any platform so you don't need to learn, test and integrate an unfamiliar load balancer if you want to migrate.

And thanks to our unique consultative approach, we can also get involved with a lot more than just load balancing issues e.g. IT infrastructure, network and third-party solutions.

Flexible pricing and licensing

We understand that in today's world your plans can change quickly. That's why we launched our Freedom License, to give our customers the flexibility to migrate across different platforms easily, quickly and without any financial penalty. Loadbalancer.org gives you the freedom to pay and deploy exactly how you want to, with clear pricing and no penalties for moving across plans or even platforms.

About the company

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 personalised support.

To discuss your load balancer requirements with load balancer experts, contact Loadbalancer.org on:



US Headquarters

Loadbalancer.org, Inc.
4550 Linden Hill Road,
Suite 201
Wilmington, DE 19808
USA

Canada

Loadbalancer.org Appliances Ltd.
300-422 Richards Street
Vancouver, BC
V6B 2Z4
Canada

EMEA Headquarters

Loadbalancer.org Ltd.
Compass House
North Harbour Business Park
Portsmouth, PO6 4PS
England, UK

Germany

Loadbalancer.org GmbH.
Tengstraße 27
80798
München
Germany