Optimising IT via Services Integration

How organisations can optimise their IT through services integration.

By John Bidgood, CTO Systal

Enterprises today need to manage a perfect balance of IT economics, application performance and security controls to meet their business application needs. This is why today's IT solutions drive greater interoperability requirements across multiple cloud and traditional data centre environments.

To achieve this, enterprises need to optimise their IT by performing services integration exercises of cloud, internet service and in-house provided solutions.

Reaching this ideal IT operating model requires more comprehensive hybrid vendor and management solutions, as well as a greater expertise in designing, building and inter-operating hybrid clouds than traditional IT services.

The best solution will ultimately allow applications and components to securely interoperate between public and private clouds and allow applications to be portable across these environments.

Over time these solutions must scale to deliver their respective benefits. Applications need to be able to move easily between environments, as where applications are hosted today might not be the best place for it tomorrow. For example, when:

- Test and Development objectives become production objectives with different service criteria.
- Public and private cloud experiences vary based on geographical and security constraints. The network experience for example will vary due to the physical Cloud location and network throughput from the end user and the network they connect from.

Application mobility must therefore be a governing principle across the hybrid cloud, whilst still dealing with the constraints of

---

running and integrating legacy applications/platforms in traditional data centres.

Building the Ideal IT Services Integration Model

The key principle to follow when building a hybrid cloud is to start by choosing the right approach for application hosting and for building the underlying physical infrastructure that supports it.

This 'application first' hosting policy is based on the economic considerations combined with the service/security constraints dictated by the business-critical applications. Using this model, we can see that:

- Private cloud is ideal for predicable workloads and custom SLAs for critical business applications, for example; data backup and internal databases. Then to plan accordingly to add resources as needed to accommodate expected growth.
- Public cloud is better where greater elasticity is needed for unpredictable workloads, for example; digital and IOT applications, where applications can be standardised to run on commoditised platforms with common SLAs.
- Traditional data centre or co-location environment for when there is no cloud migration option. i.e. when legacy IT compute and storage platforms are running key business applications.

To join these cloud and data centre environments together, an end-to-end architecture is formed based on a service catalogue of desired features, that draws upon all the IT features that will be required to host the applications. This catalogue will include all the resources in the legacy data centre, the various cloud options, the network, the security mechanisms and the digital platforms required to access the applications.

Working out which devices users connect and which digital platforms they are using determines the security segregation model and the resulting security zones that cloud will need to provide. For example, if a large number of users connect via 3rd party platforms/internet connections it may be better to treat all users as 'untrusted' to preserve security. This effectively forces all users to connect via application gateways or user VPNs instead of connecting directly to application servers.

The next consideration in the architecture is the network throughput. Network latency will typically have two considerations:

- The physical location for the hosted application services and the way network latency impacts remote users at their various geographical locations.
- Latency of servers operating within and between the cloud or data centre locations. This will include traffic between the gateway and application services as well as server replication traffic between locations.

Predictable, secure performance provided by the network is therefore essential. WAN acceleration and application load balancing can help offset some of the performance issues over distance and help manage greater levels of resilience but careful planning on placement of these devices will be needed for it to work correctly.

Buying predictable bandwidth helps, but it costs more than a 'best efforts' internet connection. So much so, that the costs of the bandwidth need to be offset by the added risk of using internet VPN wherever possible for remote offices and users. Otherwise the cost of the network bandwidth might further prohibit public cloud introduction and expansion.

Once the underlying physical infrastructure is formed, then the procurement and in-house creation of the pertinent 'as a service' models can begin. At this point the various options for ongoing deployment, management and integration as an "overlay" becomes the most challenging aspect of any hybrid cloud operation.

Managing the Services Integration Operation

To provide the optimum operational IT solution on top of the underlying infrastructure, enterprises have to find the ideal compromise between in-house and external skills to maximise the outcome of your plans and the adoption rates of cloud deployments.

Consideration also needs to be taken for when things go wrong, for example the ability to integrate these services together and perform root cause analysis when someone reports the application as 'slow.'

There is a big difference between purchasing 'Infrastructure as a Service' (IaaS) and 'Platform as a Service' (PaaS), in terms of both levels of integration and multi-site resilience. Business continuity planning is key. Many Service Managers are surprised when things go wrong and a single point of failure is identified only after the event has occurred. A traditional business recovery service might not be needed in terms of physical hosting, but the process and tools are still required to test and implement a recovery situation.

This means that enterprises should be on the lookout for potential integration partners as well as individual partners who would perform the roles traditionally filled by a Cloud Service Provider (CSP) and an Internet Service Provider (ISP). This may even extend to companies performing Security, Incident and Event Management (SIEMs) to secure the overall end-to-end service.

Ideally these potential integration partners are technology and vendor agnostic, and able to work on your behalf to secure the best long-term solutions for your business. i.e. as an IT partner their primary motivation is not to sell you hardware or bandwidth unless it's strictly necessary. Furthermore, they may offer you an opportunity to avoid margin on margin pricing by helping you to source and procure ISP and CSP services directly, and even help to manage these on your behalf.

The right optimised IT approach ultimately allows applications and components to interoperate between traditional data centres and cloud. This can be done with the freedom to dynamically provision and manage applications based on your business needs. By following the right principles early on and making careful choices to avoid vendor lock-in and nurture the right partnerships, you can scale a very cost-effective service. This can be done over many years, and prevent you undergoing a complete rejuvenation every time your network needs to change and develop.