

# Playing the SLA long game

## The role of optimised platforms

By Dale Vile, July 2012

Originally published on



<http://www.theregister.co.uk/>

When IT is your day job it's easy to lose sight of why you are actually there. Alright, from your perspective it's probably to pay the bills, fund your next holiday, buy nice stuff, and so on. But when you look at it from the point of view of whoever is paying your salary, the point is to enable and add some sort of value to the business. In effect, you, and your compatriots working in the IT department, are there to provide a service.

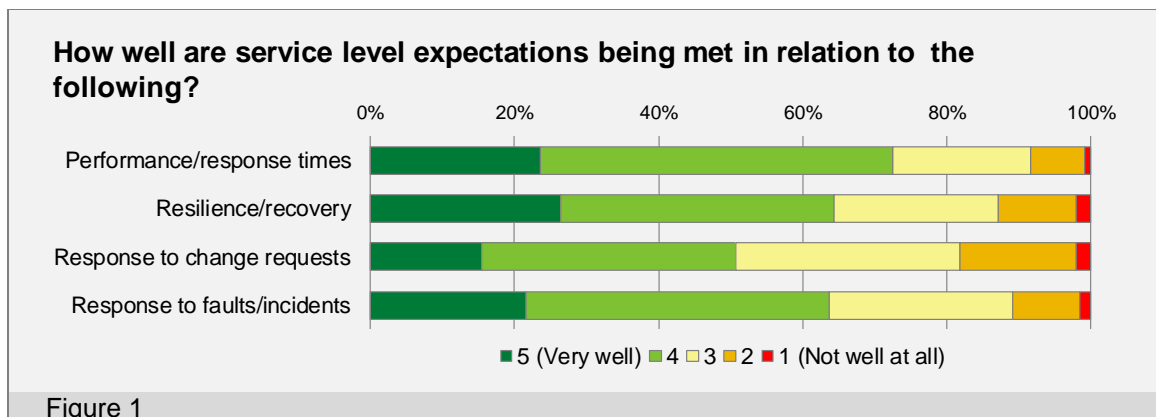
### Back to basics

How good the service is that you deliver can be measured in all kinds of ways. Apart from providing appropriate systems and capability in the first place, however, this largely boils down to a few key things:

- Maintaining system performance (response times, throughput, etc.)
- Minimising downtime and ensuring full recovery after a failure
- Responding quickly to new business requirements and change requests
- Responding quickly when unexpected problems arise

Oh yes, and you're probably also going to be judged on how efficiently you operate.

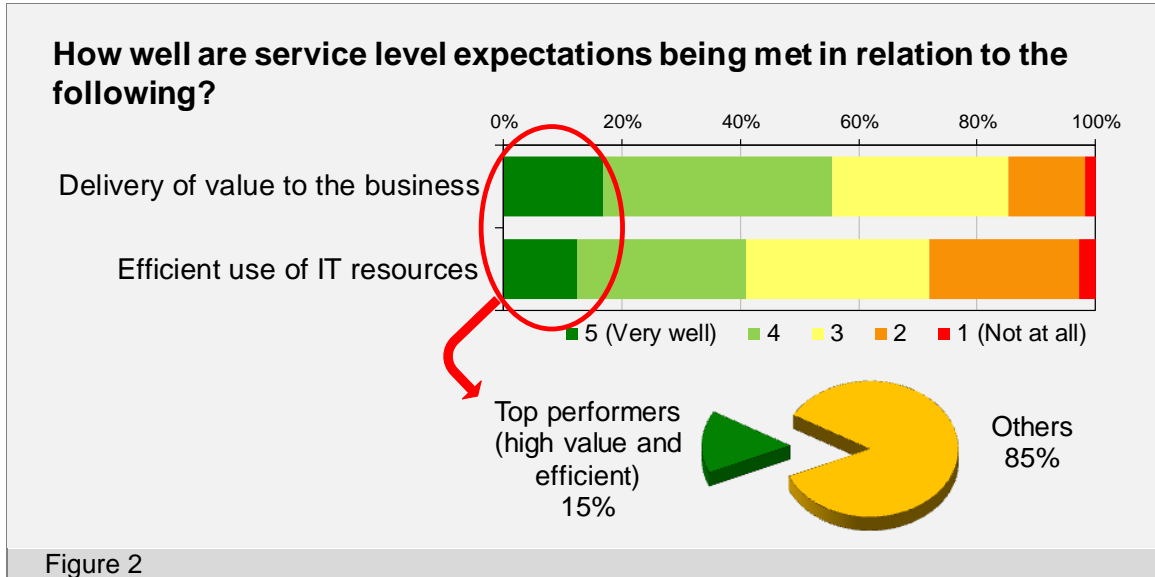
So how well are IT departments doing with all of this? The results of a recent Reg reader survey (570 respondents) threw up a broad spectrum of performance levels, so clearly some are doing a lot better than others (Figure 1).



What we are looking at here is probably a best-case representation of performance across the IT community. The survey this comes from was mostly about private cloud, which as a new(ish) type

of 'advanced' solution will have attracted a disproportionately high number of respondents from more progressive (i.e. better performing) organisations. Average performance in the population as a whole is likely to be significantly lower. Nevertheless, there is enough variability here to facilitate some useful analysis, which we'll come onto in a minute.

To help with this, we can net out the above even further, on the basis that at the simplest level, performance can be considered as a function of value and cost/efficiency. It's then possible to identify a group of particularly high performing IT departments, which represents around 15% of the survey sample (Figure 2).



So what is it that these 'top performers' are doing to achieve such dizzy heights of success when it comes to IT service delivery?

**Telling you what you already know**

When you stand back and analyse things from a cause and effect perspective (as we analysts tend to do quite a bit), you can easily come up with a list of factors that impact on IT service delivery in one way or another (Figure 3).

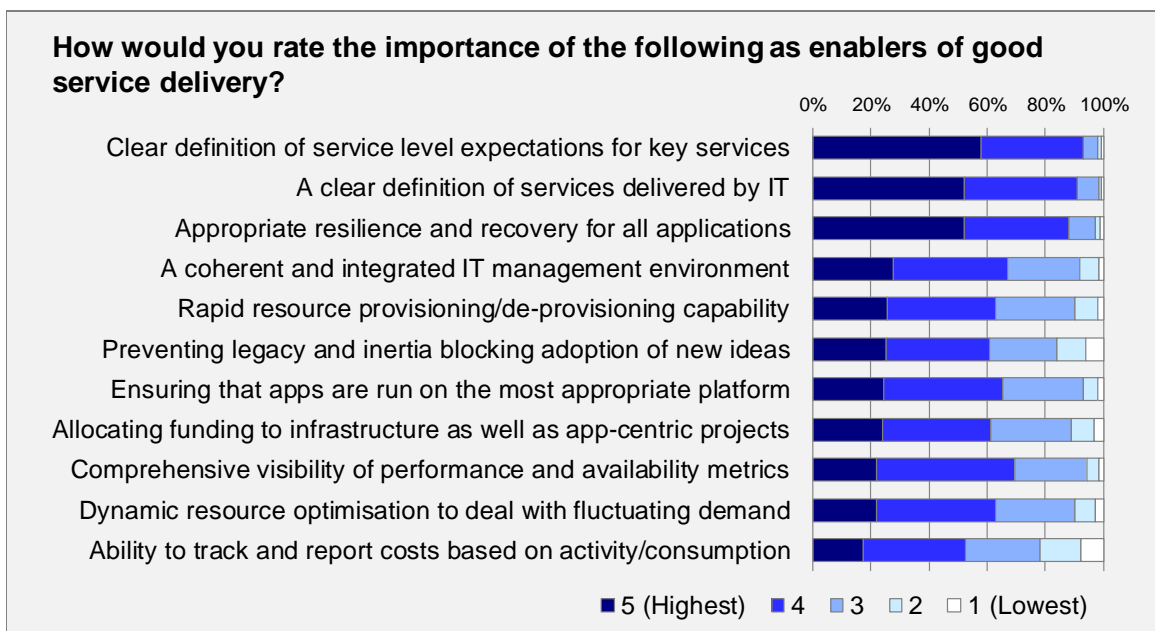


Figure 3

The fact that every item on this list was rated at an importance level of 4 or 5 out of 5 by at least 50% of respondents tells us that most people have a good feel for what they should be doing. The trouble is though, that while the top performers are delivering reasonably well against this, the majority of the 'other 85%' generally aren't (Figure 4).

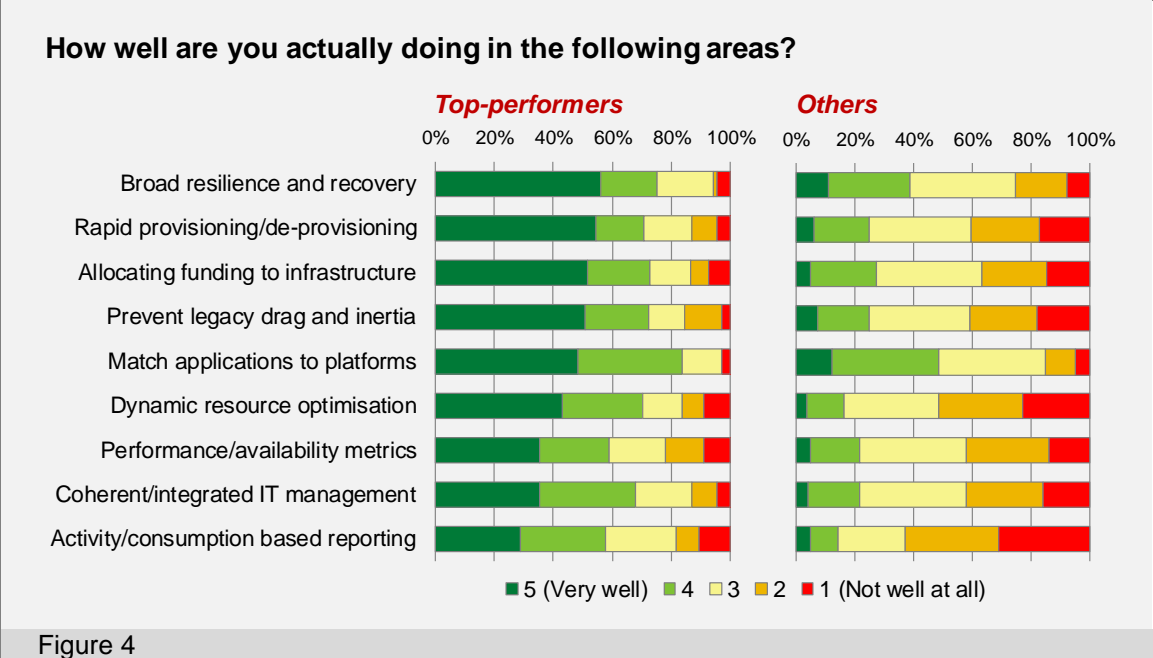


Figure 4

This picture is reminiscent of comparing report cards back at school. For some it might bring back memories of secretly hoping the smug kid who was good in class and on the playing field would someday 'get his'. Unless you were that smug kid, of course, in which case you might be tutting and shaking your head while looking across at the losers and thinking about how they can be so bad at everything.

**Loser or victim?**

But before we run away with the notion that the above picture is somehow telling us that the majority of IT departments are not doing their jobs properly, it's worth remembering some of the shortcomings are caused by factors beyond IT's control.

Take the item at the top of the list, for example - resilience and recovery. Past research studies have told us that lack of forethought or awareness on the part of those specifying requirements often means that high availability (HA) and disaster recovery (DR) requirements are not acknowledged and scoped into the plans and budgets for new systems. Even if the IT department asks the relevant questions, when funding is short, business stakeholders frequently dismiss the issue.

The problem here is that HA and DR are akin to insurance; it feels like you're spending money and getting nothing tangible in return. It's only when the first failure occurs that IT's advice is taken seriously, but by then there's no budget available to fund HA configurations or backup systems. And who gets it in neck for systems availability issues? Yep, those incompetents in IT.

**Funding hide and seek**

Lack of appropriate funding in general is responsible for a lot of the performance gaps we see. A factor that makes a huge difference to how well IT professionals can do their job is whether they are allowed to invest in infrastructure. The specific challenge here is that over the past decade, business stakeholders have increasingly looked for a direct link between money spent and short

term business value created. Budgets have therefore tended to revolve around discrete applications and systems that deliver something specific in terms of visible new capability that business users can touch and feel.

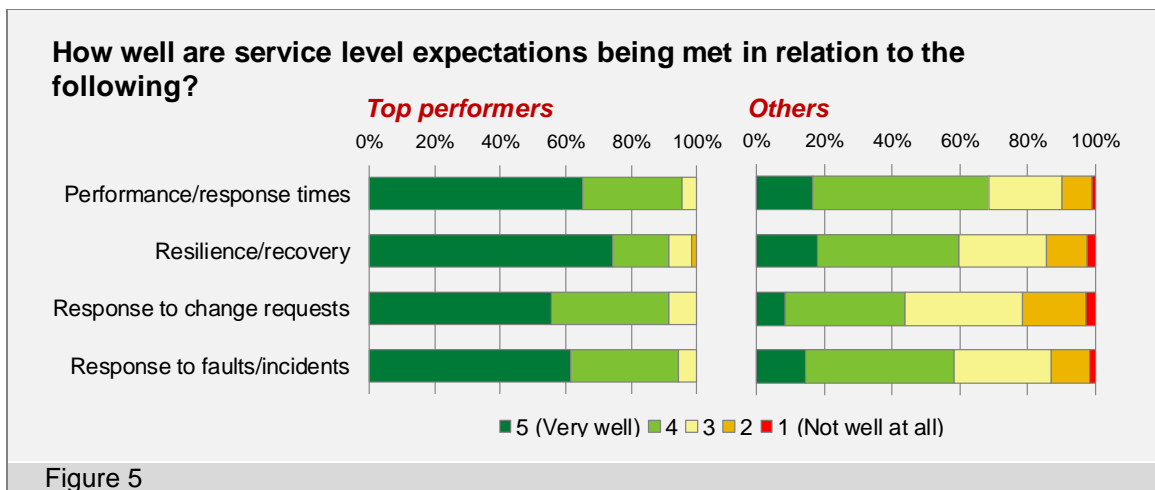
Against this background, securing funding for proactive infrastructure upgrades and future proofing measures has become more difficult. You might know that the way things are trending means your network is not going to keep up with increasing demand unless you invest in more capacity or better quality of service management capability. However, it's only when users start screaming about performance issues getting in the way of them doing their jobs that funding is (often begrudgingly) cut loose.

Then there's one of the biggest ironies that comes through in our research time and time again. Many IT teams are under constant pressure to reduce costs, but at the same time are hamstrung from an efficiency perspective by inadequate and disjointed tools that they are not allowed to upgrade or replace. Excessive reliance on error-prone manual processes that could easily be automated both perpetuates high overheads and costs, and keeps IT staff from contributing in ways that really matter.

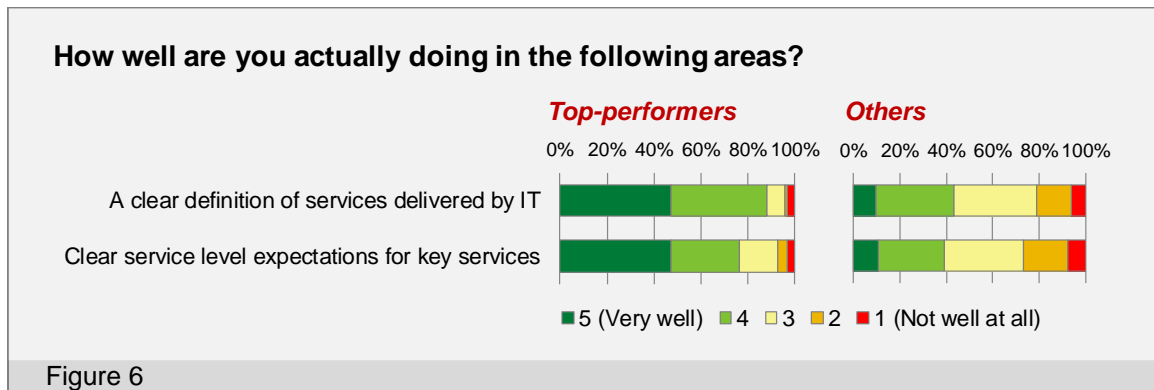
As a result of the constraints, the game being played in many IT departments is to hide infrastructure and tools investment in project budgets. Add a little bit here, a little bit there, and hope no one notices. But this only gets you so far, and coming back to our top performers, it isn't a coincidence that allocating funding to infrastructure is high on the list of enablers. That's the real way to do it.

**Knowing what matters, and focusing on it**

Whatever the reasons behind the capability gaps that exist, you can only start to improve things by establishing what really matters to the business. This brings us back to the fundamentals we were discussing right at the start. Despite the fact that the top performers are not perfect across the whole range of service enablers on our list, the overwhelming majority of them still score a 4 or 5 out of 5 against each of the fundamentals (Figure 5).



The trick to pulling this off is to know where to focus your efforts and resources. In practice this means having a clear understanding between IT and the business of the services IT is expected to deliver, and a definition of the service levels associated with the more important ones. The top performers do particularly well in this area (Figure 6).



This kind of clarity helps in two main ways. Firstly it highlights the stuff that it's critical to get right, e.g. the key services/applications and associated requirements with regard to performance, availability and so on. Secondly, it prevents IT wasting time and resources on things that aren't necessary or helpful.

**It's as much about what you don't do**

Applications vary immensely in their service level requirements. The demanding business critical systems at one end of the spectrum may genuinely need redundant high end servers, fast resilient storage, high speed networking and enterprise class middleware such as cluster-enabled application servers and database management systems, all wrapped in bullet-proof security.

At the other end of the spectrum, however, you have the application knocked up to help marketing department do the monthly analysis of partner related spend (or whatever) based on extracts from the CRM system. You might describe this as a convenience app – pretty much as far from business critical as you can get. But would the marketing guys ever admit to that when you are asking them about requirements?

In between, you have every combination of requirements in terms of performance, scalability, resilience, data protection, security, communications and access, and stakeholders, who regardless of real need, always ask for the best for their application. This can easily (and often does) lead to expensive resources (e.g. dedicated servers, premium storage, and high end database licences) being used unnecessarily, translating to wasted money on both hardware and software.

There's also the incremental operational burden. How much data do you back up nightly, for example, that hasn't changed for a week, a month or even a year? How many servers do you patch routinely that are accessed no more than once a day, once a week, or once a year and/or have single digit levels of utilisation? Meanwhile, how many genuinely critical systems sit on under-spec'd and under-protected platforms, delivering a poor level of performance that users have given up moaning about because it's been that way for so long?

You can only start to eliminate this waste and exposure by constructing some kind of service catalogue and agreeing with the business the requirements and expectations for at least the key applications and services within it. You might then need to pay attention to asset and configuration management if your tooling and processes are not up to the job, but starting with the services layer gives you more chance of drawing a line between spending money on enhanced management capability and things that matter to business people.

And lest we forget, make a point of identifying dormant applications and retire them wherever possible.

**Getting off the critical path**

Freeing up resources in the ways we have described allows you to redirect them more usefully, and a priority area here is speeding up change management. One of the most common complaints from business people, apart from IT costs too much, is IT takes too long. Systems-related work, whether

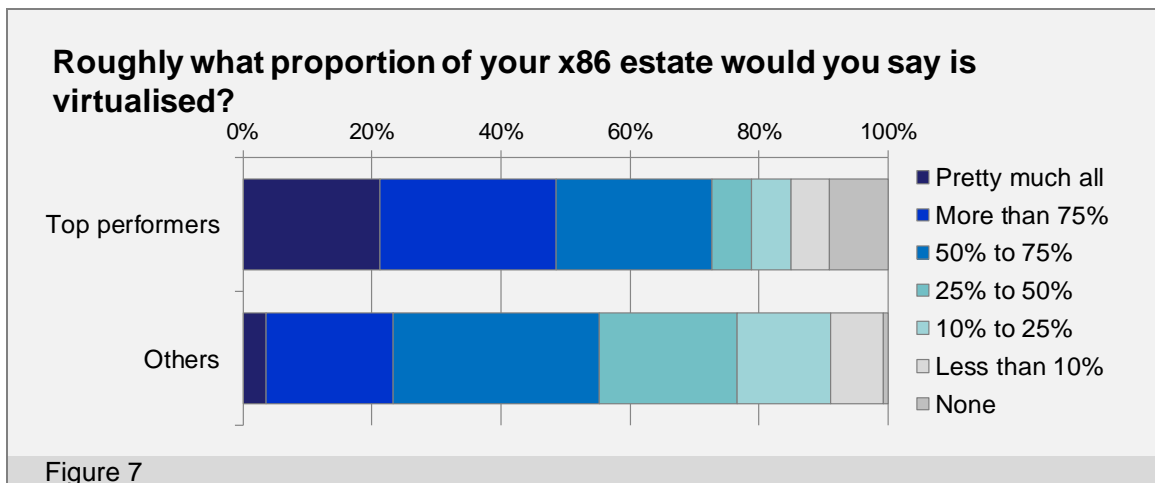
it's new functionality or changes to existing applications and services, is often on the critical path to implementing change within the business.

There's a process dimension here, and whether this can be handled through a spread sheet, home grown application or full blown software package such as a service desk and/or project portfolio management system, depends on the scale and complexity of your environment.

Introducing more responsiveness and flexibility at the platform and operations level, though, is also important. In this respect, many IT departments are discovering the potential of x86 server virtualisation, which most got into as a way of consolidating applications and other workloads onto a smaller number of physical servers. This lowered ongoing capital costs and reduced the number of physical pieces of kit that needed to be managed, which in turn meant less overhead. All very good for eliminating waste in line with our discussion above.

However, as experience with server virtualisation has been gained, more and more IT departments are appreciating the flexibility and responsiveness benefits. Being able to provision a new application via a virtual server on an existing piece of kit, for example, avoids the need to procure and configure new physical servers, which can dramatically speed up the time to deployment. Being able to migrate workloads easily between servers has then enabled a quicker response to changing demands, e.g. an application needing to scale beyond the level originally envisaged.

It's therefore no surprise to see top performers making much more use of x86 server virtualisation than other organisations (Figure 7).



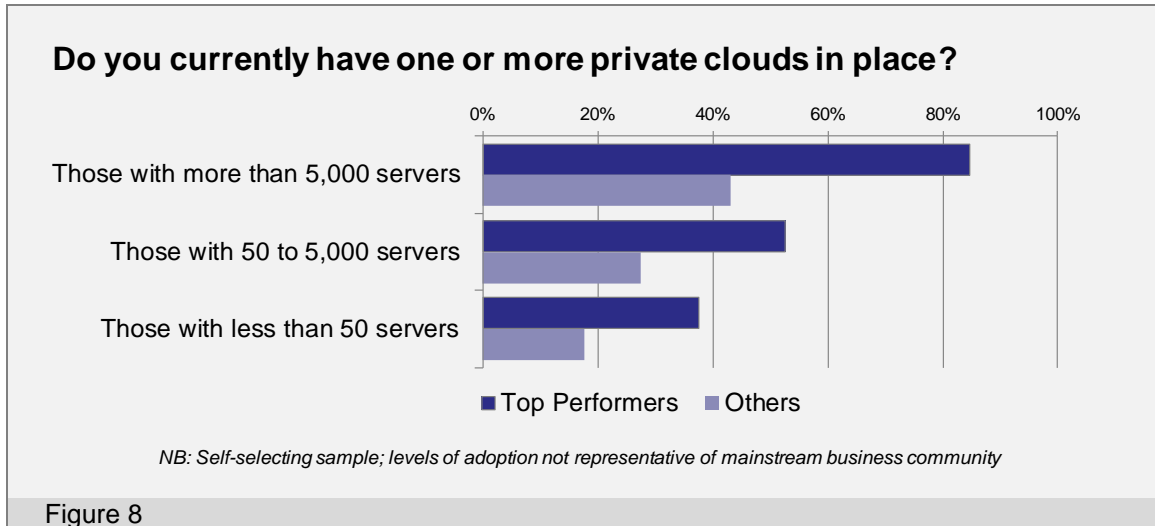
But virtualisation has not been without its issues. A number of our research studies have flagged up virtual server sprawl as becoming a real problem for many. Put simply, the ease with which virtual servers can be created mean they sometimes proliferate out of control.

Most heavy users of x86 server virtualisation have therefore had to pay attention to beefing up their management environment, and IT vendors have been working hard to deliver the capability required to meet the needs in this area. Meanwhile, some of the key vendors in the virtualisation arena have starting taking flexible x86 based platforms to the next level.

**From the 'V' word to the 'C' word**

The name of the latest game in x86 server computing is 'private cloud'. The basic idea is to pool a bunch of servers and other resources (storage and networking) to create a general purpose platform upon which a variety of workload types can be run simultaneously. An important attribute of private cloud is the rapid allocation/de-allocation of resources to/from workloads, enabling a more dynamic approach to management, and the next level in terms of service flexibility and responsiveness. Private cloud architecture can be used in your own datacentre, or set up on dedicated kit in a hosted 'co-location' style manner.

Unlike the so called 'public cloud', that is often met with a lot of scepticism and cynicism within the mainstream IT professional community, private cloud has had a more consistently positive reception. This is partly because it builds on the familiar concepts of virtualisation management, while at the same time allowing many of the benefits associated with cloud computing to be delivered under the control of the internal IT department. In our recent survey, we uncovered quite a lot of early adoption activity, particularly among top performers in organisations with large IT infrastructures (Figure 8).



Apart from the level of automation, monitoring and management, the big advantage of private cloud over more traditional virtualisation is that it is able to support demanding workloads that need to be powered by multiple servers. Through the concept of 'elasticity', resources can be allocated to and reclaimed from workloads as demands fluctuate. Together with an ability to achieve higher application availability on a broader basis, this makes private cloud great for service level management as well as rapid response to new business requirements.

### Investing for the long haul

Today more than ever, platforms, tools and techniques are available that allow a step change in the efficiency and effectiveness of IT service level delivery. While we haven't been exhaustive here, we have touched on some of the more important principles and developments.

But, and this is a very big 'BUT' - none of this counts for anything unless business and IT stakeholders jointly recognise the importance of prioritising enhancements to the way in which IT services are delivered, and provide the funding and resources necessary to invest for the future. So, it's over to you.

## About Freeform Dynamics



Freeform Dynamics is a research and analysis firm. We track and report on the business impact of developments in the IT and communications sectors.

As part of this, we use an innovative research methodology to gather feedback directly from those involved in IT strategy, planning, procurement and implementation. Our output is therefore grounded in real-world practicality for use by mainstream business and IT professionals.

For further information or to subscribe to the Freeform Dynamics free research service, please visit [www.freeformdynamics.com](http://www.freeformdynamics.com) or contact us via [info@freeformdynamics.com](mailto:info@freeformdynamics.com).

### Terms of Use

This document is Copyright 2012 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process. Hosting of the document for download on the Web and/or mass distribution of the document by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd.

This document is provided for your general information and use only. Neither Freeform Dynamics Ltd nor any third parties provide any warranty or guarantee as to the suitability of the information provided within it for any particular purpose.