

IT infrastructure for mid-sized organizations

If you want to get there, start here

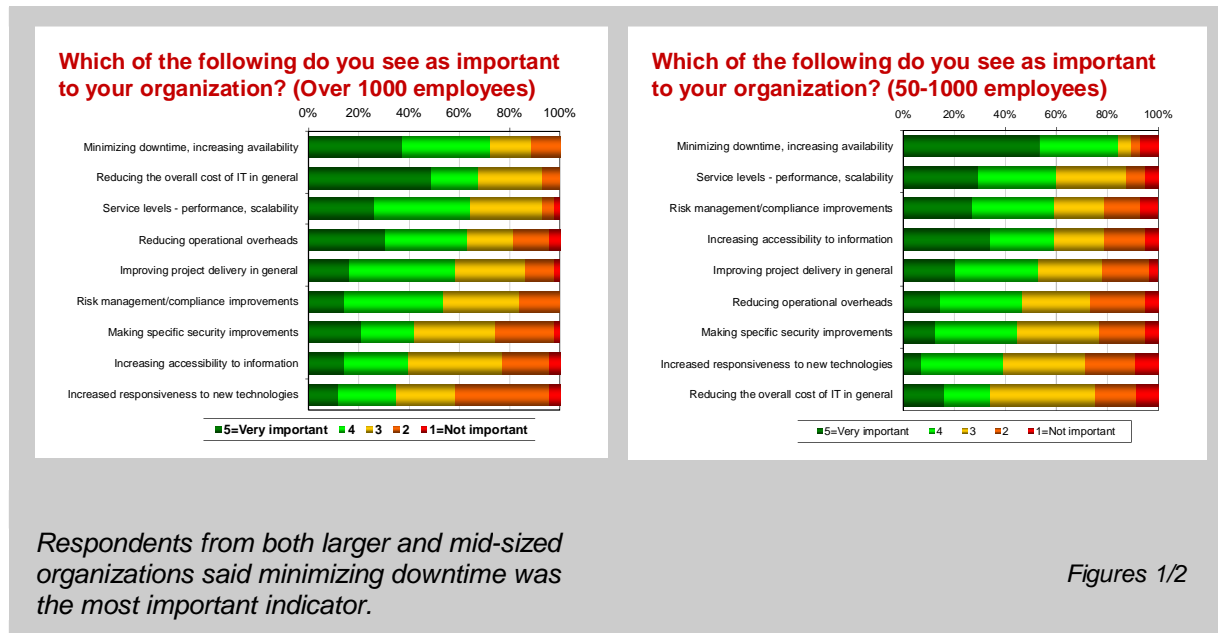
Jon Collins, Freeform Dynamics, June 2010

IT infrastructure is often presented in a way that is more relevant to the largest organizations with the biggest wallets. But if your business isn't operating at 'enterprise scale', how can you think about infrastructure in order to make the most of new developments and build a platform for the future?

What's it like to be a mid-sized organization?

In the IT industry, the "mid-market" tends to be defined as the set of organizations that don't fit into a clearly defined category: organizations that don't qualify as either "enterprise" (large-scale multinationals that operate as nation-states), nor as small businesses with no more than tens of staff. This category is frequently underserved, particularly when IT infrastructure is discussed.

In this digest, based on a workshop conducted in partnership with The Register, we review what matters and what doesn't to non-enterprise companies as they look at moving their IT infrastructures forward. To kick off the workshop we conducted a short survey (128 respondents) to gain a picture of such organizations. The survey was aimed particularly at understanding where best practice and infrastructure technology investment can best deliver on what organizations see as important (Figures 1 and 2).



This report is based on the outputs of an online workshop and research study completed in January 2010. The research involved gathering feedback from 128 predominantly IT professionals. The report was sponsored by Oracle, though the study was designed, executed, analyzed and interpreted on a completely independent basis by Freeform Dynamics.



As illustrated in the figures, the number one priority for mid-sized organizations was to minimize downtime and increase availability - in fact, this came out as the number one criterion for organizations across the board. What's interesting is how cost comes out so low for businesses with 50-1000 employees. Of course, cost will always be important – but this perhaps reflects the fact that such organizations are focusing more on keeping IT running. For larger companies, cost comes out as the number two criterion, which may be due to IT departments in such organizations needing to respond to cost-cutting edicts coming down from board level.

With such factors in mind, the aim of this digest is to offer advice and guidance on how best to look at IT infrastructure as a whole, and how to take things forward in a way that balances what is already in place, with meeting the evolving needs of the business. This digest summarizes the workshop articles, which covered:

- Understanding the services you are trying to deliver
- Building a picture of the IT you have in place, and how it is structured
- Defining how you manage your information assets
- Balancing trade-offs when it comes to investing in new technology
- Managing the IT environment, both as it stands and into the future

If you're in a company that doesn't fit the "enterprise" mould, if you're bored of hearing what enterprises are doing with IT and you want something more relevant to your own needs, read on.

Starting with service delivery

Infrastructure exists to help process and manage the information required by the business, to support how it delivers services to its customers. In this first section, we consider IT from the business perspective, and look at how a service-based view can help. We also consider how information needs to be managed, and how such aspects as compliance and risk can be taken into account.

A good starting point for companies of any size is to remember IT is a means to an end – and the end is helping the business deliver services to its customers. IT helps support this by processing and managing the information required by the business.

But how does this translate into service delivery terms? We all know whether or not we are getting what we would perceive as 'good service', be it in a restaurant, a shop or indeed from IT. The challenge is knowing what to put in place to support the delivery of services, whatever their shape or form.

While all organizations are different, a simple way to think about different service requirements from an IT perspective is to remember that business activities tend to fall into one of two categories:

- Activities that are, or can benefit from being executed in a structured manner;
- And those that cannot.

Let's consider the more structured activities first. These have been given various names at various points over the past few decades – workflows, business processes and so on. Methodologies abound to identify and restructure them, with the goal of improving their efficiency and repeatability. Such techniques can also serve as the basis for understanding what can be automated, and indeed, service expectations from the business perspective.

For IT to support structured activities, it needs to operate efficiently, that is, as cheaply as possible. Structured interactions are largely transactional, in that they respond to specific service requests or updating specific information items – consider an accounts receivable process, for example, or a customer support workflow. In the structured world, if a service is unavailable or inaccessible for whatever reason, the cost to the business can be very high. As a result, flexibility is a second priority to service levels.

The second activity type – "everything that cannot be structured" – is no less important when it comes to deciding IT service requirements. For less structured activities, accessibility, flexibility and

responsiveness are key drivers. There is no point at all in having a world-class customer management system, for example, if it isn't able to provide a phone number to the service engineer who is working off-site.

Before you embark on making any changes to your infrastructure, a key question to be answered is, "What services does your organization want to provide, and therefore, what requirements does IT need to meet?" For example, your systems may need to support high volumes of transactions, or perform to specific availability criteria. The service delivery characteristics you define will then dictate the kind of IT environment you need.

Knowing what you have

Having spoken to senior decision makers in companies of all sizes, we know that having a clear understanding of the existing context and environment is an essential pre-requisite for moving forward. This is also seen as an architectural view – but 'architecture' can be overloaded with complexity and meaning which is often more relevant to larger companies. In this section, therefore, we consider how to build a view of what exists, and what is required from the business perspective to enable this.

No organization has the luxury of designing its ideal IT environment from scratch. Having spoken to senior decision makers in companies of all sizes, we know that to maintain or improve service delivery, you will need a handle on what is already in place. In practical terms this equates to understanding:

- What IT systems and resources you have, both physical and logical
- What information repositories you have
- And what dependencies exist between all three – systems, services and information

This may sound obvious, but we are frequently reminded from our research activities that few organizations ever have a comprehensive understanding of the above. Fortunately, attaining this understanding need not be a very onerous task, particularly if you start with the Pareto principle in mind: can you identify the 20% of your IT facilities that deliver 80% of the value to your business?

Key to this kind of exercise is the 'means to an end' factor. If you haven't already documented what you have, we're pretty confident that just doing so will turn up some things which merit attention. The bigger-picture view you gain will also provide a starting point for defining an IT architecture that fits with the needs of the services you want to deliver.

It's important not to get bogged down with terms like IT architecture, however. Documented approaches tend almost immediately to big frameworks and complex, boil-the-ocean models that are only really applicable to larger organizations (and then, probably, only in moderation). But there is a great deal of merit in the idea of designing your IT environment both as a whole, and as individual parts.

From some of our previous [research](#) [1], we know that poorly architected IT environments tend to suffer from more downtime than those for which resilience has been built in from the start. They will also be more costly to administer and deliver lower levels of service. IT architecture is not the solution to all these ills – just having a well structured environment will only get you so far. It does, however, offer the starting point for doing a better job, more efficiently.

Indeed, if you work out all the things you are trying to do and then map them against what you have in place, you may find that you want to keep things as they are. Equally, however, you may spot some incremental improvements you can make to your existing environment; or indeed, you may decide more drastic action is required. In all cases, at least you will be starting from a clearer point than "let's just do it and see what happens" – an approach which most organizations have been guilty of, at some point.

Managing your information assets

The architecture of the IT environment needs to encompass both the way application workloads run on servers, the desktop estate and access mechanisms, and the way information is managed, secured and protected. In this section we drill into this latter aspect – storage – to illustrate some of the issues involved.

A number of challenges are associated with managing information today, not least data growth, which remains a primary driver for how IT infrastructure in general, and storage in particular is delivered. In addition, data duplication across the infrastructure, and the resulting fragmentation of information assets, place an additional burden on the storage environment, not to mention the people and processes needing to work with it.

Against this background, it is useful to consider first the ‘outside-in’ view of what storage needs to provide:

- Firstly, its role is to deliver data to applications and users consistently and efficiently: that is, as and when needed, at the required levels of throughput performance, and at an appropriate cost.
- Second, storage should also be able to recover from failure situations. It is one thing when things are going right; quite another if things go wrong. Here we can think about backup and recovery as well as the ability to replicate between storage arrays, and indeed across sites.
- And finally, storage management needs to consider the operational constraints of the storage environment itself. This is not just about having visibility on what storage exists, but also to respond to changing conditions and changing requirements, preferably as automatically as possible.

Responding to these needs requires more than buying in just a bunch of disks. For example, storage arrays tend to fall into one of two categories – “high-end” for very expensive, high-performance disks, and “mid-range” for the rest (you don’t hear of “low-end” storage arrays, but you do hear of “entry level”). Deciding what data should go where can be quite a skill, particularly as data characteristics change over time. A well-managed tiered storage set up would match the ratio of high and lower cost storage in use, with the requirements of the data at any given time.

An additional dimension is to meet requirements on data availability. These are generally achieved by placing a copy of the information in some other (preferably safe) place – either online (for example, via a second array to which all data is replicated) or offline (for example, on tape, in a fire-proof safe or at an off-site storage facility). However the plethora of options in this area – from backups and archiving, to snapshots, replication and so on can just make things more complicated.

Where to start? As discussed with IT architecture in general, if you can build a clear picture of your storage requirements (remembering not all information is created equal – the 80/20 rule can be used here), together with a map of what capabilities you already have in place, you’re already half-way there.

From here, the next step is to produce the map of how you’d like your storage capabilities to look, based on your information needs and your broader infrastructure plans, for example vis-à-vis any intentions you might have around virtualization. Given the plethora of options, each of which comes at a cost, storage planning will always be a compromise between functionality and affordability.

As a result it is worth thinking about how storage technologies might be used in tandem. For example, while it still isn’t cheap (though it’s getting cheaper), implementing de-duplication might provide immediate savings in terms of bandwidth and latency reduction. However, by considering its impact a little more broadly, the bandwidth savings may now allow (for example) data replication to another site, making disaster recovery possible whereas in the past it was not.

Storage then, is like a team of Sherpas; it does the heavy lifting so the rest of IT can make its way up the mountain without having to worry about the provisions. But it needs to be designed for the long haul, if it is to deliver on the value it promises.

Making the right bets

Each option comes at a cost, so storage planning will always be a compromise between functionality and affordability – which brings us to a discussion about the economic aspects of IT. This section considers the idea of ‘value’ and its measurement; it also considers keeping a picture of costs following deployment, and how to consider IT as a portfolio of assets which can be managed both individually and as a whole.

IT budgets have never been that easy to come by in mid-sized organizations, and the current economic climate is not helping. Making investments in IT can sometimes seem like a betting game – given multiple possible ways to respond to a given need (be it to grow the business or make things more efficient), how can organizations decide between the different options and approach IT investments in a cost-effective manner?

The importance of a solid business case cannot be underestimated, but equally, it does not have to be an inches-thick document that covers every aspect. It should be able to answer some pretty basic questions, however:

- Do we really need it?
- What else has been considered and ruled out?
- What do we stand to gain from the acquisition, and by how much?
- How can success be maximized?
- What issues are going to prevent things from working, and how can they be avoided?

Much of the business case hinges on the term “value”, which can be summed up by the simple equation:

$$\text{Value} = \text{Benefits} - \text{Costs}$$

This may sound obvious, in hindsight – but a good understanding of value also depends on having a clear view of the business benefits of the selected options, together with a comprehensive picture of all the costs involved.

Not everything will even reach the business case stage. Running IT is a constant case of juggling priorities and dealing with more urgent requirements, often at the detriment of what might be more important. It is no coincidence that much IT marketing follows the “who shouts loudest” model, in the hope that certain agendas will be pushed up the priority stack and perhaps get signed off.

Several factors mean that financial prudence doesn’t end on the day after deployment. IT has a natural lifetime, beyond which service calls start to go up, and the costs of maintenance begin to exceed the benefits achieved. Of course, it would be good to spot when this point is being approached, but often it is easier to plan the replacement/upgrade in advance, so that new systems can be in place before the edge of the cliff is reached.

Fundamentally, keeping the IT infrastructure going is little different to managing a pool of assets. Whether a certain element of IT exists in-house or is bought in as a service, whether it is operated by employees or contractors, whether it is funded in a single payment that must then be amortized, or whether it is leased over time, all these things will make a difference to the overall cost of each service in the portfolio.

The “advanced class” of IT economics considers IT services not just individually, but in terms of the way the whole portfolio brings value to the business. While this brings things back to the balance sheet, managing the portfolio is just as much about understanding the relevance of each IT service to the business need, as its tangible costs. The portfolio needs to be kept up to date, and it is important to work at the level of individual assets particularly at their creation, but also when the time comes to retire them. To be blunt, if a service is neither relevant nor efficient, what is it still doing there?

Post-deployment best practice

The rhetoric and frameworks of IT management best practice are often seen as more applicable to enterprise organizations. In this section we consider the best practice state of play when it comes to systems and service monitoring and management, and ask the question: how does “good enough” apply to operational best practice for mid-sized organizations?

Much best practice assumes that the operational management of IT follows some kind of maturity model. Many organizations are, in theory, chaotic; a logical next step is to implement a layer of control, and from there one might say one could start building a managed environment. After that it is reputedly straightforward to achieve fully dynamic IT management.

The problem with this model is that it simply isn't true. The 'chaos' myth – that a majority of IT departments are doing no more than fire fighting, shoring up the defenses and keeping the lights on – doesn't actually represent the majority of IT organizations out there. In reality, most organizations are working at a level we could term “good enough”: the last time we researched this [directly](#) [2] we found the challenge was more one of running through the mud, than fighting fires.

In truth, most mid-sized organizations hover somewhere between a 'controlled' and a 'managed' level of IT management capability. When, on occasion, things tend towards chaos, for example after a security breach or a more significant systems failure, budget can miraculously become available to sort things out. When it comes to financing more general improvements however, systems management tools and techniques have not been seen as a major priority. “Good enough” is a phrase that can cut both ways.

The 'mid-market' has traditionally been under-served by management software vendors, who have put their efforts into providing big-ticket enterprise management offerings. Until quite recently something of a gap has existed between the simpler management tooling suitable for use in smaller companies and the management suites deployed by the largest organizations.

All the same, there is still scope for improvement in non-enterprise companies. The primary goal is one of efficiency, which doesn't have to be about tools at all. A training course on writing administration scripts, for example, could have a very quick payback. And there is much to be said for implementing the right policies and processes before looking at tooling. Experienced practitioners understand that most systems management projects are exercises in people and process management, just as much as dealing with technology challenges.

It may be that some recent developments in IT infrastructure – notably virtualization, which encourages the more flexible use of physical resources – will in the future require IT management to adopt more dynamic approaches to IT management. In the meantime, and for the majority of organizations who just want 'good enough' IT, any such ambitions may be a step too far. “Good enough” IT requires putting in place good enough management tools and practices, which aim to address the challenges of delivering a managed IT environment. While this may offer a potential stepping stone for more dynamic approaches in the future, for many organizations this will be sufficient as an end in itself.

Building a sustainable platform

Throughout this workshop, a continuous thread has been about getting the balance right – in terms of architecture, investment and best practice. But if IT can never be finished, what can organizations do to ensure that the balance is maintained into the future? As well as the broader question of sustainability, this section also considers the more specific factor of 'green IT' – is it possible to consider what is, in essence, a wide array of manufactured goods without being hypocritical about being 'green'?

As a starting point, let's recall that the essence of best practice is about what IT enables, rather than what IT is. Whichever statistics you choose to quote (for example Gartner's Simon Mingay), IT only constitutes a few percentage points of an organization's total carbon footprint. Of course, there will always be merit in trying to reduce this figure in absolute terms, either for altruistic or business value

terms – there is no point in wasting valuable resources unnecessarily. But this line of thinking is not going to get you much closer to the answer.

Meanwhile, we have “what IT enables”. The fact is that IT has a long history of under-achieving, for a whole number of reasons. As we have already mentioned in this paper, the theory was (and indeed remains) that IT would help drive super-efficient, hyper-dynamic businesses. However the reality for many organizations is that the business trudges on, inefficient as ever, with an overlay of automation that can seem more about keeping the corporate lights on than adding any inherent business value.

With or without the green imperative, IT has a massive, still largely untapped role to play in making the business run far more efficiently. If we do think about this in carbon terms, this equates to the 98% of corporate emissions. But it also equates to a substantial proportion of financial costs, service metrics and performance indicators, all of which can have scope for significant improvement.

To bring things full circle, this boils down to good governance, such that IT service delivery is appropriately mapped to the needs and priorities of the business. While, as we have illustrated, this does require some thinking about, it doesn't become any less important. IT is powerful stuff, but as long as it is implemented reactively, it will only have a limited impact. Meanwhile, of course, if it is used in such a way as to improve, change or even eliminate processes, and thus their carbon impact, there's plenty to get excited about.

From our research we know that the opportunity exists to deliver IT in the right manner, in a way that it actually achieves its objectives – be it in terms of unifying business communications, helping decision makers access the information they need, consolidating infrastructure and rendering it more responsive to changing needs through virtualization, better management tools and so on.

This goes right to the heart of what we mean by sustainability. Some organizations might be doomed to spend eternity on the back foot, running inefficiently and falling foul of whatever legislation comes their way. But for others, who are not afraid to take the initiative and look at how IT can really make them more efficient and effective, there is all to play for.

Taking things forward

Where to go from here? The IT landscape continues to evolve, and a number of recent advances in infrastructure such as virtualization, software as a service and pooled storage do indeed have benefits for non-enterprise organizations. However, such technology developments tend to have a horizontal impact: to be implemented successfully, they need to be considered in terms of the whole infrastructure, not just one element.

While cost reduction may not be the main driver for mid-sized organizations, the current financial climate nonetheless offers the opportunity to stop and think not only about how IT is being used across the board, but also what potential exists to improve how services are delivered to the business – either by exploiting new technologies, or just by making better use of what is already in place.

Whatever initiatives or plans you currently have underway, we hope you have found this digest useful. If you have any feedback, let us know at info@freeformdynamics.com.

References

- | | |
|---|--------------|
| [1] Risk and Resilience | June 2008 |
| [2] Relieving the Systems Management Burden | January 2008 |

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 ITC strategy, planning, procurement and implementation. Our output is therefore grounded in real-world practicality for use by mainstream IT and business professionals.

For further information or to subscribe to the Freeform Dynamics free research service, please visit www.freeformdynamics.com or contact us via info@freeformdynamics.com.

About Oracle



Oracle (NASDAQ: ORCL) is the world's largest business software company. For more information about Oracle, please visit our web site at <http://www.oracle.com>.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Terms of Use

This report is Copyright 2010 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.

This report 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.