

A Smart Guide published by



Evolution of desktop service delivery

Aligning client computing and virtualisation technologies with business needs



With Compliments



Technology needs context

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

An innovative research methodology allows us to gather feedback directly from those involved in IT strategy, planning, procurement and implementation.

Our output is grounded in real-world practicality for use by mainstream business and IT professionals.

For further information or to subscribe to the Freeform Dynamics research service, visit our website or contact us at

info@freeformdynamics.com



freeformdynamics.com

Evolution of desktop service delivery

Aligning client computing and virtualisation technologies with business needs

Terms of Use

This document 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. 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.

Contents

The genie is out of the bottle	4
Why a client computing strategy?	6
Starting with the business and its users	9
Guiding your client computing strategy	14
Drilling into specifics	20
Typically managed desktop/laptop computing model	22
Session virtualisation	22
Virtual Desktop Infrastructure (VDI)	23
Application virtualisation	24
Client-hosted virtual machine	25
Personal computer alternatives – tablets, pads and slates	26
Making the right choices	26
Moving the strategy forward	29

The genie is out of the bottle

Few would deny we are at a tipping point in terms of how we interact with computers. The range of device options is diversifying: every week, technology sites tell us of the latest form factors and clever gizmos, and how they are set to transform our personal and working lives. Meanwhile in the data centre, server computers are powerful enough to run multiple 'virtual' machines, supporting new and evolving ways to deliver IT services to devices of all types. The umbrella terms 'desktop virtualisation' and 'cloud computing' are indicative of some of the new models possible.

Working practices are evolving in parallel, as broadband networking and mobile technologies make it possible to work from almost anywhere, and with anyone. Some companies no longer have pre-allocated desk space, preferring 'hot desking' facilities for when people come into the office. Benefits of home working are well-documented, from saving commute times (estimates of about 60 hours average per year have been quoted) to increasing both work done and employee well-being.

Such developments can result in a number of pressures on the IT department, which has to cope with the widening set of technology choices and working models. If left unchecked, IT running costs can increase dramatically, for example, in terms of operational management, user support and data integration. Meanwhile, risks to productivity, security, licensing, and power management also rise.

However, many IT decision makers are finding they have little choice but to respond to their users' changing practices. While the business cannot afford a chaotic IT environment in which anything goes, it is clearly becoming important to think more broadly than standard approaches where applications are deployed to desktops and laptops alone.

NOTE: We use the term 'consumerisation' to describe how users are using home computers for work, say, or relying on non-corporate web sites to collaborate with their peers. While Freeform Dynamics research suggests some organisations prefer 'locked down' IT environments, we know from experience that attempts to restrict this phenomenon can result in a backlash, hidden from the eyes of those trying to control it.

As a starting point for success we need to consider all compute devices in the round – that is, desktop and laptop computers, thin clients, smart-phones and Personal Digital Assistants (PDAs), notebooks and slate-type devices – as well as the variety of ways that such devices access centrally provided IT services. We shall refer to all types of user-facing device as 'clients', and their selection and use to access centrally provided IT services as 'client computing'.

Given the wealth of choices and combinations, is it any wonder that it is so difficult to know how to decide on an appropriate strategy? Rest assured it is not our intention to lead you towards any one solution or another, as each business is different.

In this guide we provide a practical framework for client computing that puts the needs of your business and its users first, without undermining your ability to provide a cost-effective IT service. It covers what matters in terms of guiding principles, capabilities and potential benefits, and then explains how to approach the needs of your own organisation in a way that stands the test of time, not just by keeping an eye on future device types, but also by ensuring investment in new technology is protected.

If you're a business or IT decision maker with an interest in how to exploit the potential of client computing to further the success of your business, read on.

Why a client computing strategy?

Before continuing, we need to mention that few organisations today are looking to undertake a complete replacement of their desktop and laptop estates with some esoteric alternative. However, based on the organisations that we speak to, we do see an increasing number of usage scenarios that need to be addressed. You may find any of the following familiar:

- Your support desks are faced with an increasing volume of calls about device types that are not part of the 'corporate standard', or usage models that were not envisaged, possibly with activity in locations that are difficult to reach. The result is that issues are harder to diagnose, and managing user perception of IT support is becoming more difficult.

- A number of executives are requesting non-standard equipment, operating systems and applications, and cannot understand why it should be a problem, despite such purchases falling outside current arrangements in terms of licensing, procurement and so on.
- Someone at a senior management level recently read an article on new desktop computing models such as desktop virtualisation, and asked you to investigate whether your own organisation could benefit.
- You are conscious that a proportion of your user base requires a relatively fixed set of applications and services, and you are keen to explore what you see as a range of potential alternatives for delivering capability to this group.
- Your existing environment is feeling the strain of both new applications and raised expectations, for example, use of unified communications, rich internet applications, video training or desktop sharing, which business users at all levels feel should 'just work'.
- You are looking to update your current client computing environment as part of a regular upgrade and replacement cycle, and you simply want to be sure that you are not missing a trick when it comes to alternative options and that your investment will last.

If such scenarios are treated in an uncontrolled manner, the end result is piecemeal solutions and, indeed, fragmented management. The goal of having a client computing strategy is

not only to meet the needs of your user base, but also to enable you to recentralise the management of all types of client and application.

As we have already highlighted, the keyword for any such strategy is ‘balance’ – of course you will want to provide your users with the best possible environment, but you will also want to do so affordably and in a way that fits with the objectives of the business. Equally, you will want the confidence that whatever you put in place today will remain appropriate for the duration of the investment.

NOTE: The greatest level of stability and future proofing is achieved when the strategy and approach adopted deals with fundamental business requirements, but is also tolerant of changes in more volatile areas such as devices, specific business practices, regulation, and so on.

So, where to start? Despite first impressions of weekly changes in working practices, the fundamentals of underlying business requirements do not move that quickly. Neither should they: environments that change all the time are bad for users, who end up spending more time learning how to get on with new technologies, than being productive in their jobs.

Many initiatives fail not because they were technically inadequate, but because they didn’t take the needs of the organisation and its users into account. Rather than looking for quick fixes or betting on any one option in the hope that it will deliver, the only sensible approach is to put the business first.

Starting with the business and its users

While the expression ‘people are an organisation’s greatest asset’ may be over-used, the two are interlinked: profitable organisations and productive people go hand-in-hand. The starting point for client computing is to consider that what is good for business users is also going to be good for the organisation as a whole. However, this does not mean a free-for-all. Business management is about harnessing the talents and experience of staff and providing appropriate tools to support their activities. In turn, the client computing strategy needs to take into account the careful balance between making the most of an organisation’s people, whilst at the same time delivering on the goals of the business.

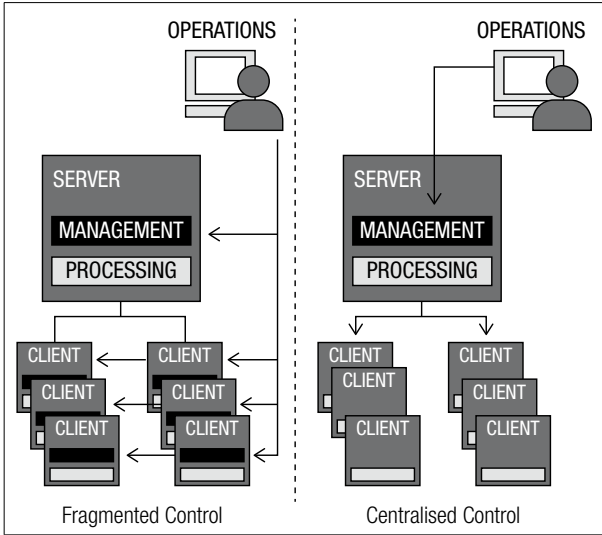
Given the fact that technology will continue to evolve, we need to look beyond specific technology options to deliver on this aspiration. Our years of experience looking at technology use in business suggest a number of principles that will remain constant even as technology changes, not least:

No single technology is, or ever will be ideal for all circumstances. Some technologies are more of a commodity, offering a lower cost way of meeting the majority of the organisation’s needs in a given area. Others will be more appropriate for specific requirements, potentially at a higher cost. Therefore, the IT environment will continue to require approaches that take into account how to integrate and manage a variety of complementary equipment types.

Devices will continue to play a fundamental role in the user experience. The user interface of the device, and the applications and services it makes available, are critical to the effectiveness of individual users. This will inevitably drive diversity in the device pool as groups of users work with device and service combinations that best meet their needs.

Processing happens somewhere, with a power implication. A misunderstanding associated with some client device models is that they are lower-powered and therefore more efficient and/or better for the environment. From a technical perspective, if processing doesn't happen on the client, then it will need to take place on a server, which will require additional cooling. This is a trade-off: sometimes it makes more sense to run a smaller client and do server-side processing, whilst in other cases it will be better to make use of client-side processing.

Business-critical applications and data should be centrally controlled and managed. This remains true whether applications execute on servers or clients, whether the services are being hosted 'in the cloud' or on-premise, and through whichever device they are surfaced. Only with centralised control and policy setting can the IT function offer a cost-effective service to the business which keeps control of risks. This also enables trade-offs to be made in terms of where an application executes, depending on network bandwidth and latency, performance, security and related cost constraints.



Moving from fragmented to centralised management

Individuals will continue to ‘do their own thing’ with available technologies. Consumerisation works both ways, either using corporate equipment for personal needs, or using personal kit for business. Policy setting and education therefore starts with the user rather than with the device, even if certain policies can be automated through device features – for example, PIN setting on mobile phones.

Only a subset of applications and services are in fact business critical. Effective IT management processes take into account the business criticality of applications and services. As the environment changes it is important to monitor new tools and

usage patterns, such that IT can 'adopt' them as business critical if appropriate. For example, an organisation may decide that the 'LinkedIn' web site is business critical, whereas the 'MySpace' web site is not.

Security and continuity risks will evolve in line with technology and working practices. IT faces the law of unintended consequences – as new models become prevalent, so the 'bad guys' will devise ways to exploit them. This is not only about cybercriminals: most security breaches are internal, and will more likely be caused by ignorance than malice. Business continuity also needs to extend in line with how people work – data is just as subject to flood, fire or theft in the home environment as in the office.

As well as these fundamentals, a client computing strategy needs to deliver against measurable benefits whilst minimising costs and risks. The right combination of client devices can drive business value and productivity in a number of beneficial ways, including the ability to:

- Respond to business events, whether involving a structured business process, or working in collaborative situations, or reacting to unexpected circumstances.
- Create, modify and share information in an accurate, appropriate and secure fashion – by information we include text and structured data, video and other media.
- Access the information and services required to do the job, either in isolation or working with others.

- Work effectively and consistently from a variety of locations and across time zones, both individually and with others, whether connected to a network or not.

Client capabilities need to be delivered at an acceptable cost. There's more to consider here than the initial purchase price – and unforeseen expenses can nullify any benefits if not considered comprehensively. Specific costs include:

- Acquisition – not just book price but also the time required to define, agree and procure the chosen solution. An individual spending a day deciding on a new mobile phone is not a good use of their time.
- Infrastructure – a delivery model may have an impact on the back-end IT environment, for example, requiring server storage, networking or existing applications to be upgraded or replaced.
- Integration – for example, to enable information to be accessed in the right form on different device types.
- Management and support – help-desk operatives can be ill-equipped to respond to queries, for example, about how to connect devices to the corporate network. New delivery models have an inevitable impact on management processes and tools, which may need updating.
- Information protection and security – for example, defining appropriate mechanisms to encrypt or back-up information stored on a device.

- Licensing – including overheads incurred in determining what hardware and software assets exist, whether they are covered by existing licensing agreements, or whether these vary by client device.

Such cost criteria remain applicable, whether for a single device or for a major procurement. In some cases the issue might appear trivial – the ability to display email correctly for example – but the impact can be profound if the device or service is seen as business critical.

We can't cover every eventuality – if compromises and trade-offs need to be made, they need to be the right ones. We consider these in the next section.

Guiding your client computing strategy

As we have already discussed, while technology options will continue to evolve, the fundamentals of business change relatively slowly. Of course it is important to evaluate business requirements first, as to do otherwise would put the cart before the horse. This approach also has the advantage that a client computing strategy can be founded on a more stable base than technology alone.

The best way to understand what the business needs from its client computing strategy is to profile its users. In this exercise you learn not just who is using what applications, but also which working practices and usage models are the norm. The Pareto principle (80/20 rule) very much applies – the goal is to end up

with a simple set of user categories, against which technology solutions can be mapped.

NOTE: While user profiling does not need to be onerous (for smaller organisations, all you really need to kick things off is a clipboard and pen), it can become quite daunting either because of the scale of the organisation, or because of the number of user groups and applications in use. Automatic discovery tools may help identify applications, or you may choose to limit the exercise to one part of the organisation or user group – home workers, say.

In all cases, a key element is to get out there and talk to real people about their working practices. It is dangerous to second guess how people might work – small nuances, for example, misunderstandings about how sales people log their client visits, may have a big impact on success.

A profiling exercise enables you to identify a number of categories of user. There is no hard and fast rule as to the categories you might end up with, but certain criteria dictate how the boundaries fall. For example:

- **Performance.** The specific applications and services required, and the platform requirements of the applications – for example, whether they are compute- and/or graphics-intensive. Note that even ‘simple’, often unavoidable activities such as web browsing can require quite significant processing, particularly if rich content such as Flash animation is involved.
- **User mobility.** The locations and travel habits of the user group – for example, whether they spend their time at home or

on the road, the connection types available, and the activities they undertake in different places.

- **Flexibility.** Whether the required set of applications is relatively static or whether it is likely to change, for example, based on changes to working practices or frequent upgrades to the latest software releases.
- **Collaboration.** The level to which individuals need to work together, using what kinds of tools (for example, unified communications and/or video conferencing), to achieve what ends and within what kinds of constraints, such as team-working across time zones.
- **Infrastructure.** The availability of supporting infrastructure, for example, reliable networking versus largely offline access, quantities of data to be transferred between users or with internal systems, server processing required and so on.
- **Security.** Security requirements relating to the nature of the data being handled by the user, as well as risks associated with the users and their roles.
- **Control.** The degree to which applications need to be centrally managed. This will depend on the level of technical competence of the users, their administrative rights, and how much support they need.

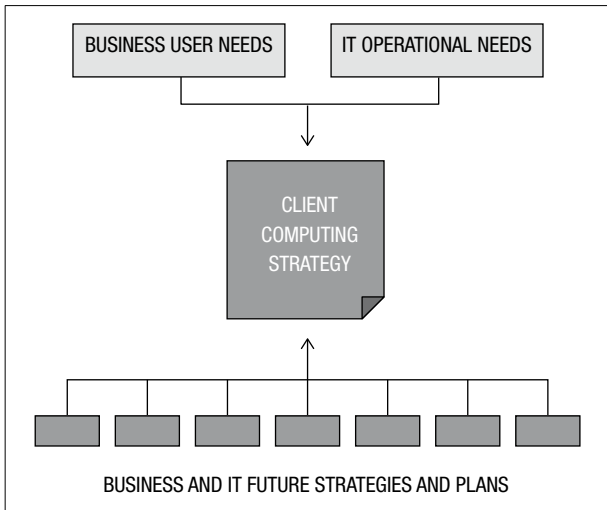
By working through criteria such as these, you can identify user categories that are appropriate to your organisation. The table below should give you a start in terms of what you can look for, and the kinds of categories that make sense.

User category	Examples and criteria
IT development staff	Software developers, architects and testers requiring access to a changing set of applications, but based in an office
IT operations/support staff	Technical experts and administrative staff using a core set of applications, but who will on occasion require access to a broader application set
General professional users	Senior management and administrative business users making relatively light and predictable use of email, office tools and so on
Home workers	Like general professional workers, requiring an equal level of collaboration, but spending most of their time at home with occasional visits to the office
Transaction workers	Business users in fixed roles, mostly using enterprise applications in a routine, prescriptive manner
Other (non-IT) technical and creative staff	R&D staff, engineers, technical designers/architects requiring access to more graphics intensive applications
Office-based power users	Business users with complex application needs, for example, finance staff, marketing teams and business analysts
Mobile and remote professional users	Similar to power users, but also requiring access when on the road or from remote locations, for example, sales executives, roaming managers and service engineers

No set of categories can be perfect: you need to decide what works best for you. As a result of your discussions with users and their management you will also gain useful insight about what's planned, to help ensure the future safety of any approach you select. Finally, you will almost certainly end up with a handful of specific individuals that don't seem to fit into any group – while you may need to treat these as exceptions, at least you can also get on with treating the norm.

Armed with this information, you can start to draw up a plan. Some categories may have given clear signals that they are struggling with their current set-up, while others may appear to have everything they need. You can use such information as the basis of prioritisation – for example, you may have defined a 'road warriors' category that is under-served, and therefore should be considered for treatment first.

NOTE: At this point you may be able to identify potential 'quick wins' – areas where a solution may give an immediate return, for example, to resolve a specific issue. Quick wins serve to get more general buy-in at an early stage of any initiative.



Balancing user and business needs with IT needs

Before looking at options, remember that a client computing strategy does not exist in isolation from other areas of IT, nor indeed, business strategy. It is therefore important to have conversations with relevant business and IT stakeholders, to gain insight into areas such as:

- Applications, both packaged and bespoke
- Data management, compliance and governance
- Management and support both of centralised equipment, and of users and devices
- Data centre and network architecture

- Power monitoring and management
- Mobility and communications
- Sourcing and service provision
- Security, and the broader areas of compliance and risk

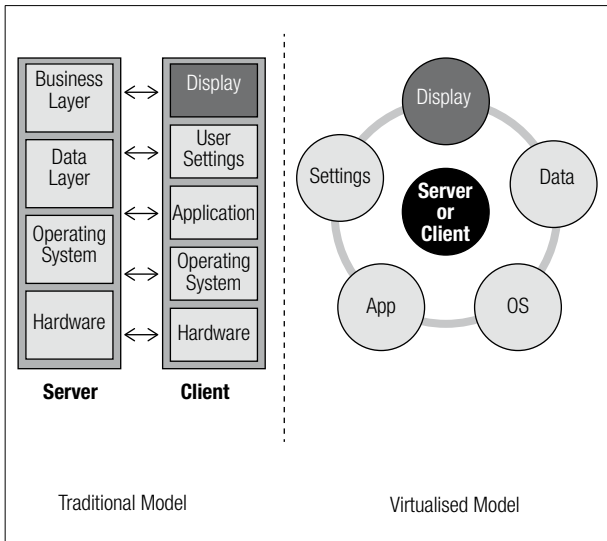
Such conversations will reveal plans that could have a significant impact on the client computing strategy. Keeping the future needs and strategic direction of the business in mind will help to avoid making decisions that may quickly be superseded, or indeed investing in technologies that could quickly become inadequate for the tasks at hand.

Drilling into specifics

Where to start with the nitty-gritty of technology selection?

The next stage involves mapping the benefits of specific client computing technologies against the categories of users and their requirements, all the while keeping an eye on the manageability of the resulting environment.

To be clear, IT has moved a long way from client-versus-server computing. In traditional IT models, applications have run as a 'stack' on the desktop PC, with a back-end application stack on the server. Virtualisation on both server and client means that anything can run anywhere – as illustrated in the figure below.



Virtualisation breaks the bounds of the traditional IT stack

We have already mentioned the importance of running a centrally managed IT organisation – but the options are broad in terms of what needs to run where. Each configuration has trade-offs, for example, in terms of ensuring sufficient processing capability is available on the server or the client, or appropriate networking capability is in place. In this section we run through some of the options that exist today and what they bring to the party, before mapping their benefits against the user categories we presented before.

Typically managed desktop/laptop computing model

The client platform most familiar to users and IT departments currently is based around the personal computer, either in its desktop or laptop form. This needs no introduction – but it is important to note how generations of PC have improved in terms of performance and power efficiency, making certain things simple on the desktop today (for example, the use of video, voice or complex data manipulation) that were difficult or expensive to achieve before. In general Microsoft Windows is the dominant operating system, but Apple Macs are prevalent in the design community and among fashion-conscious executives, while Linux has achieved in-roads particularly in more technical environments.

Benefits of this model include flexibility, based on the fact that suitably configured modern desktops and laptops are adequate for the majority of applications. Meanwhile however, power management can be a challenge particularly for older models. Manageability is sometimes cited as an issue, for example, in terms of operating system patching and keeping end point security mechanisms (such as anti-virus and personal firewalls) up-to-date – though newer generations of PC can incorporate capabilities to help address issues of remote management and security.

Session virtualisation

In this model (traditionally known as ‘thin client’), multiple desktop clients access applications from an instance of a server-based operating system. In essence, single applications or the entire desktop run on a server, and display graphics are transmitted to

the desktop client. The client can be 'thin' in that it only needs to process and display graphics – but it can equally be a window on a standard PC, or indeed in a web browser on any other device. Meanwhile 'ultra-thin' or 'zero' clients contain simply a network and a monitor port, relying on specialised hardware to decode the remote session.

Primary benefits of this model are centralisation of control and data storage. It can also be quite flexible by offering multiple configurations for individual user categories, but the approach depends on a network being available and so is more suited to corporate office environments. Defining a session virtualisation architecture relies on having a good grasp of exactly what applications need to run, particularly for communications tools and graphics- or compute-intensive applications, as these will consume server CPU cycles.

Virtual Desktop Infrastructure (VDI)

From an architectural perspective VDI follows a similar principle as 'session virtualisation' in that core processing takes place on the server, with the display transmitted to the client. In this model however, each user has their own operating system on the server, running as a virtual machine isolated from other users.

As with session virtualisation, the main benefit of VDI is centralised management, in particular for operating system images. Equally the model provides flexibility in offering different application sets for specific user categories. Given the use of virtual machines, it also has availability benefits since a virtual machine can more easily be copied or moved than a physical machine.

Again like session virtualisation, VDI requires that the memory and processing requirements of each virtual machine are taken into account (including graphics processing) and needs a persistent network connection of sufficient bandwidth. Given that each user is allocated his/her own virtual machine, VDI also demands more server and storage resources than session virtualisation. Note that recent offerings permit the virtual machine to move locally to the client for off-line operation (see client-hosted virtual machine, below).

Application virtualisation

Traditionally, desktop applications have required installation on the client, sharing (sometimes conflicting) resources such as run-time libraries. In the application virtualisation model, individual applications are bundled with the resources they need to run and can then be streamed to the client device as and when they are needed. Applications are isolated from other software, removing the potential for conflicts.

The main benefit of application virtualisation is user flexibility, as applications can be linked to users rather than specific hardware. It also offers application stability and reliability, given that every application has its own 'protected' configuration, without the performance hit that comes from isolating complete virtual machines. From a management perspective, IT managers can build, deploy and centrally manage cleaner operating system images and reduce the time they spend on interoperability testing, application deployment and patch management – though this

does incur an up-front overhead. Once the application 'bundle' has been streamed to the client it can be executed whether or not the client is connected to the network, catering to mobility needs. Application virtualisation can also reduce licensing costs by providing access to certain applications only when needed.

Client-hosted virtual machine

In this model a virtual machine is run directly on the local PC, more usually within an environment (known as a 'Type 2 hypervisor') running on top of the operating system. A more recent development is that the virtual machine environment runs directly on the machine's hardware – this is known as the 'Type 1 hypervisor' model.

The single-desktop approach offers the flexibility for individual users to keep multiple client environments apart from each other, for example, to separate business from non-core working. The approach is also popular for software developers, to create target configurations for test purposes. The model has management benefits, in that the virtual machine image can be controlled and distributed centrally, for example, to pre-configure a virtual machine to assure compatibility with a legacy application. When combined with VDI, the single desktop virtual machine approach can improve the loading of users per server. However it does increase complexity by introducing new layers, and therefore potential security risks.

Personal computer alternatives – tablets, pads and slates

Lastly, it is worth mentioning the newest kids on the block, in terms of fashion anyway. The principle behind the 'slate' form factor has been around for decades, and while Microsoft first released its tablet PC operating system back in 2001, the model failed to ignite the market until quite recently. All this looks set to change however, with the introduction of products such as Apple's iPad, similar products running Google Android, and slicker tablet PCs from the Microsoft stable.

A significant benefit of these form factors is portability, sitting between fully fledged computers and PDAs, with sufficient processing power and a decent enough screen size to allow a broad range of tasks. Downsides from a corporate perspective include security and manageability, particularly with newer operating systems which do not have as comprehensive a set of supporting tools.

Making the right choices

We clearly haven't covered all the options here – we could have mentioned net books for example, OS streaming, or other virtualisation models which allow for a desktop environment to be carried around on a USB stick. We could also have drilled into PDAs and smart phones, and combinations such as running session virtualisation on a slate. Equally, some options can be used in combination – application virtualisation can be used in conjunction with VDI, or session virtualisation with laptop PCs.

The point of this guide is not to be exhaustive – rather, it is to illustrate how each option brings with it a range of benefits which can be mapped against user categories. As new technologies and combinations become prevalent, you should be able to evaluate the benefits they bring and map them in a way that makes sense to your organisation. Looking back at the list of criteria from the previous section, we illustrate some of the mappings in the table below – note that your own specific category definitions may lead to different mappings.

	Desktop/laptop PC	Session virtualisation	Application virtualisation	Virtual desktop infrastructure	Client-hosted virtual machine	PC alternatives (tablets/pads/slates)
Performance	●	◐	◐	◐	●	◐
User mobility	●	○	◐	○	◐	●
Flexibility	●	○	●	◐	●	●
Collaboration	●	◐	●	◐	●	◐
Infrastructure	○	●	◐	●	●	●
Security	◐	◐	◐	◐	◐	◐
Control	◐	●	◐	●	◐	○

Key: ● Fully applicable ◐ Partially applicable ○ Not applicable

It is worth working through some of the user categories and considering how their requirements might be met:

General professional users. The standard desktop model will already be the norm for this group of light-footprint users, and there will undoubtedly be a view of 'if it ain't broke, don't fix it'. There may be management benefits to be had from looking at a VDI model for this group, however these benefits need to be weighed up against the loss of flexibility, for example, in terms of allowing for more powerful applications in the future.

Remote professional users. This group may need a combination of options, starting with a reasonably specified laptop PC for use at the home office or on the road, and potentially supplemented by other models to augment the base level of functionality. For example, you could deploy application virtualisation to ensure that this group has access to non-standard applications when they need them.

Transaction workers. This group, using a more limited set of applications (most likely browser or terminal based), may be candidates for the session virtualisation model, once again with the caveat that you should consider what the application set will look like two to three years from now – for example, whether unified communications capabilities will need to be incorporated.

As you can probably see, the key to selecting the right mix of client models for your organisation rests with the criteria you define. Get these right and you end up with a straightforward set of user groups, as well as a clear picture of appropriate options.

The less desirable alternative is to play ‘what if’ games, throwing new client types at users and hoping they will be acceptable.

Moving the strategy forward

We hope that, in reading this guide, you can see how to identify client computing models that are appropriate to your own organisation.

NOTE: At this point you may have identified multiple user categories that could benefit from the same approach, in which case it makes sense to treat them together.

Following this comes the tougher stuff, namely planning how to take the chosen options forward, from defining a business case through to deployment. To finish, we would like to leave you with some tips and tricks we have learned from our research.

Set a realistic timeframe. To fit with the wider goals and budget cycles of the organisation, the strategy should support the creation of a practical roadmap to cover 12 months, together with selected options that will deliver over a three year timeframe. Anything shorter will incur too great an overhead in terms of integration, education and so on.

Look for incremental change. While the IT industry makes everything sound revolutionary, in fact smaller changes, augmenting what you have in place, are often the more cost-effective. There will be an even greater case for this as more devices flood the market – trying to keep up is not an option.

Architect for success. We know from the organisations we research that the success of most client computing approaches is dependent on getting the server and storage-side architecture right. If in doubt, run a pilot so you can test the performance characteristics of the model based on the planned application set.

Aim for managed service delivery. The problem with unmanaged client models is the unexpected cost when things start to go wrong. Whatever the approach, users will still expect the help-desk to support their own devices, and data protection challenges only get worse if left unchecked. So it is essential to build in management from the start on both client and server side, for example, in terms of image and asset management of virtual machine instances.

The device is the service. For many users, the equipment you provide is seen as the full package. So all parts – operating system, applications, back-end integration and indeed the client device – need to work end-to-end. This will remain true even if specific devices and/or services are provisioned and managed by a third-party service provider.

Think about tomorrow's needs as well as today's. Given how quickly things are changing, the worst thing you could do is put in place a solution which is unable to support a key capability just months later. Review any plans from a what-if perspective, and be sure you speak to relevant parties so you know what is coming. If you deploy too low-spec an environment, it may need replacing sooner than expected.

Let your users do their own research. While IT managers might watch with alarm as users try out new devices, as long as the risks are understood it is worth offering users the opportunity to test alternative models and see what works for them, for example, by setting up an open pilot environment and soliciting feedback.

Know when to bring on-board new technology. You will no doubt know when any new technology moves from optional to a business-critical state, for example, due to the amount of support calls you receive. When this happens, you should be prepared to work with the business units concerned, setting appropriate criteria to enable you to manage and support effectively.

Separate hype from reality, and watch out for the sales pitch. The premise behind this guide is that things are changing fast in terms of what's available – the downside is that some offerings are still relatively new, and early in their development. So, while you can expect solutions to improve over time, prepare to do your own due diligence. In addition, don't fall for vendor claims that a solution works for a broad segment of users when in fact its application is quite niche, or that it will work for future changes that haven't been fully tested (for example, moving to new browser versions).

There will never again be a one-size-fits-all model. Your resulting environment may well be a hybrid of traditional and new options, and there is nothing wrong with that – nor is it wrong to simply stick with what you have. We are moving to a world in which the number of options is only going to increase, and some

of your users may end up with multiple devices, each meeting a particular need. The important thing is to ensure you are covered from a management and centralised control perspective, whatever the combination.

Finally, and most importantly, always remember to put the business first!

We hope you have found this guide useful, and we wish you every success in the development of your client computing strategy.

This is an exciting time for anyone involved in setting strategy for client computing. The rapid innovation in this area is affecting us at Intel as it is in all other enterprise-class companies. We're subject to a wide range of new technology innovations – like enhanced device security and client virtualisation, both client-hosted and server-hosted – which influences the working practices of Intel employees, the way that IT delivers services, and the variety of devices we are able to support.

Our job is to make the best use of the available technologies to drive up employees' productivity while maintaining the costs of running IT. By understanding the needs of our users and the contexts within which they work, we can pick the right device, the right client architecture, and the right application deployment solution.

The coming years will no doubt bring us all kinds of new technology configurations and combinations. While it's difficult to know exactly what the future holds, this guide should give you a head start in deciding what's best for your business. Good luck!

Diane M. Bryant,
Chief Information Officer (CIO), Intel Corporation

This book was written independently by Freeform Dynamics and sponsored by Intel



Faced with an increasing range of computer device types and application delivery options, coupled with changing working practices and the growing tendency for staff to make use of their own kit, IT decision makers can find it difficult to define how best to manage and deliver desktop services whilst ensuring the needs of business users are met and risk to the business is minimised. Not all options are equal so it is important to take a 'business first' approach, which balances user needs with the cost/benefits associated with the technologies involved.

This guide provides a practical decision-making framework to assess the client computing needs of your own organisation. It covers the trade-offs, benefits and costs of different options available today, based on a set of principles which should remain valid even as the range of options continues to evolve.

If you are pro-actively developing a client computing strategy, or you simply feel it is time to review what is available and need help in getting your thoughts straight, then this guide is for you.