FUJITSU

White paper
Virtual Client Computing

Virtual Client Computing helps improve service quality and security, increase flexibility, and reduce costs. Fujitsu provides Virtual Client Computing solutions based on best-in-class virtualization technologies, proven infrastructure products, and end-to-end services from a single source. Customers benefit from rapid implementation and reduced risk resulting from Fujitsu's extensive project experience.

Contents

Challenges with PC workplaces

During the past few decades, PC workplaces have become indispensable for end users. They represent an important productivity tool supporting end users in doing their business.

Each end user may have an individualized environment to fit personal needs. Traditional PC workplaces are fully isolated from each other, i.e. you won't suffer from the failure of any other PC nearby. Besides online usage, PC workplaces can be used offline, even if there is no network available. Whenever mobility is required, people will use notebooks or other mobile devices, which in turn enables a flexible way of working.

Standard software is basically available for all purposes and can be run without any adaptation. On traditional PC workplaces, applications run locally, and user data are stored locally as well. Results of applications are directly displayed on the attached monitor. PC workplaces feature high-quality graphics, video, sound and an easy integration of peripherals (e.g. printers and scanners).

PC workplaces are easy to use, thus keeping the amount of education and training of end users at a low level. The result of all this is an excellent user experience, which is the reason that PC workplaces are well established in businesses and proven.

However, we should not ignore the fact, that PC workplaces also raise a number of serious challenges, especially for the IT organization. Lifecycle management has proved to become increasingly complex. The deployment of new PC workplaces and the application compatibility tests to identify application conflicts may end up in a cumbersome and time-consuming task. Moreover, most of the employees transform their PC they initially received with a standard user environment into one which is unique to them. And the more personal the PC workplace, the more difficult it is to manage. There are various user types, such as task workers, knowledge workers, power users, external users and mobile workers who all have different requirements and therefore have to be treated differently. PC hardware is distributed all over the enterprise, there are more and more remote and mobile users who are not always connected, but have to receive software updates and patches on a regular basis. And above all, due to the tight coupling of hardware, operating system and applications, any change in one of these layers will impact the neighboring layers, too.

When a PC fails, it will usually take a considerable amount of time until the problem is solved and the PC is available again for productive usage. Associated with troubleshooting activities, sometimes even expensive desk-side visits will occur. And of course, the impact of PC downtime is only extrapolated when it comes to a disaster on a broader scale. As a backup by all end users on a regular basis is not ensured, there is always the risk of data loss. Further security risks result from the multitude of vulnerabilities which can potentially cause unauthorized manipulation or destruction of data, as well as data theft. All these risks conflict drastically with regulatory compliance.

Another issue with PC workplaces is under-utilization of resources. People speak about less than 1% of CPU utilization on average, considering a 24 hours day, what sounds rather inefficient. However, due to the PC hardware being distributed, there is hardly any resource pooling option. No matter whether a PC is highly utilized or not, it will always consume energy while being switched on. And some people wonder, whether it is efficient to store the operating system and all the applications on the hard disk of each PC.

It should also be noted that you may access your workplace environment only, if you are where your PC is, or if you take your PC with you. Besides, there is no flexible access option.

For all reasons mentioned, the total cost of PC workplaces is extremely high and mostly not fully transparent.

Rethink your workplace architecture

As mentioned earlier, the main cause for the high complexity of traditional PC workplaces is the tight coupling between its individual layers, which makes the layers and their lifecycles dependent from each other. Breaking this tight coupling correspondingly will lead to independency of these layers and their lifecycles, will reduce complexity, simplify management, increase flexibility and speed, and finally reduce costs. Separating neighboring layers from each other is exactly what virtualization stands for.

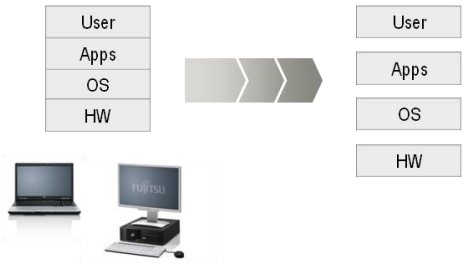

Virtualization can happen between any of the layers, between hardware and operating system, between operating system and applications, and you can even separate user personality from your workplace image. The more layers you separate from each other, the more advantages you can achieve, of course.

Server based computing

Basically, virtualization of IT workplaces is nothing new. Missing IT cost transparency and security concerns with traditional PC workplaces were the main drivers for server-based computing in the past. In a traditional server-based computing scenario, also denoted as "Hosted Shared Desktop", applications run on a central terminal server farm; and user data is located on central data storage, too. Access devices are needed for interaction only. Hence instead of a fully equipped PC workplace, a thin client is sufficient, whose function is limited to managing mouse and keyboard inputs, and refreshing the attached display with screen updates coming from the server. The two most important protocols for the communication between clients and servers are RDP (Remote Desktop Protocol) from Microsoft and ICA (Independent Computing Architecture) from Citrix whose further development is named HDX (High Definition User Experience). For standard applications which typically have only low bandwidth requirements, both protocols enable excellent remote access even over telephone lines.

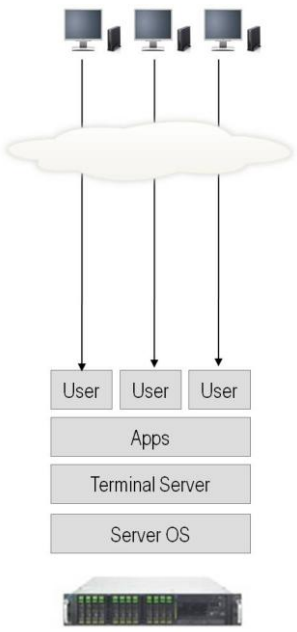

When a user wants to use an application, the thin client sends a request to the terminal server. The load balancer, an important component of a terminal server environment, checks which server has got sufficient resources; then the client is connected to that server. On the server, a new session is started, and the user is logged onto this session with his profile. The user's profile and data directories are dynamically delivered through a request to the Active Directory.

The most well-known examples of Terminal Server solutions are Citrix XenApp and Microsoft RDS (Remote Desktop Solution).

As terminal server infrastructures are centrally located, they can be centrally managed. Software can easily be deployed and updated without touching the numerous clients. Whenever hardware or software bugs have to be fixed, there is no need for IT staff to travel. Moreover, synergies can be used for managing servers and desktops. Backup takes place centrally and in a simplified manner, without depending on any end user. Furthermore, server-based computing eliminates the danger that arises, when PC users try to act as their own administrators and install software that could be a risk for system and data.

Companies that suffered from the disaster of the 11th September 2001 stated that recovery of their servers took three days, while recovery of their desktops took two years. With server-based computing, disaster recovery concepts for servers may be one-to-one applied to desktops. Besides that, hardware and software upgrades are fully transparent to the end user. Thereby, you will tremendously increase your level of desktop application availability.

As new software or growing computing needs do not require PC upgrades any more, the lifetime of the client hardware is considerably extended. Thin clients have fewer moving parts, and they operate without failure for at least twice as long as normal PCs. In case a thin client device once breaks down, it can simply be replaced by a new one. Furthermore, the usage of thin clients holds enormous energy cost savings, even when considering the required infrastructure in the data center.

The risk of data theft is eliminated, because all data is hosted centrally in the data center. In addition, the RDP and ICA protocol support encryption for data transfer and Web access, which improves security as well.

Certainly, server-based computing improves resource utilization, because resources can be shared by many users. Load balancing mechanisms take care of acceptable response times. Applications are no longer installed on a high number of client systems, which leads to a considerable reduction regarding overall disk storage space.

It is irrelevant from which device you will access your desktop environment in the data center, applications and data are accessible from anywhere. I.e. the user profile is a roaming profile and follows the user instead of staying on the client. This gives you a high degree of flexibility.

Among all client computing concepts, server-based computing provides the greatest ratio of users-to-system for the lowest TCO. According to various studies, more than 90% of Fortune 1,000 companies have deployed server-based computing, but they use it across less than 10% of the installed desktops. This underlines that server-based computing is a well-accepted concept. But there are evidently a couple of things that have kept customers from introducing server-based computing concepts despite all their advantages mentioned earlier.

Although local peripherals are theoretically supported, smooth operation is not always guaranteed, especially when it comes to the simultaneous operation of several peripherals attached to one thin client, or to synchronizing a mobile device with your office environment. By the same token, in order to make use of printing functions on a big scale, additional solutions are necessary. In particular with low network bandwidths between clients and servers, certain applications, e.g. multi-media and graphics applications, will not run satisfactorily without special solutions. In addition, the idea of mobility is not supported.

Doubtlessly, workplace management is highly simplified by server-based computing; however it is worth mentioning that even thin clients require some minor management activities, which represents an additional task for the IT department.

One of the most essential barriers that have kept users from server-based computing is the fact that desktops and applications are shared by several users. Problems posed by one user can have an impact on other users, i.e. there is no isolation between different users. Likewise, users don't get their dedicated personalized desktops as they do in conjunction with a traditional PC workplace. The changed user experience creates resistance against introducing server-based computing concepts. Furthermore, a mandatory prerequisite for server-based computing is the multi-user capability of applications. Hence certain applications have to undergo a particular treatment before installing them on a terminal server. Achieving this, can be quite a complex task.

Exactly these drawbacks are addressed by Virtual Desktop Infrastructure (VDI), a further development of server-based computing.

Virtual Desktop Infrastructure (VDI)

In a Virtual Desktop Infrastructure, also denoted as "Hosted Virtual Desktop", individual desktops with different types and versions of operating systems run as virtual machines on servers in the data center. Desktops are no longer shared; they are isolated and therefore fully protected from each other. As traditional PCs, they are individual to fit personal needs. And in contrast to traditional server-based computing, applications need not be adapted. Computing load is scaled across multiple CPUs and servers.

Virtual desktops can be flexibly accessed anywhere and anytime from any access device. For the remote access via a remote display protocol, e.g. RDP or ICA, a thin client is sufficient.

As for server virtualization, a hypervisor makes hardware and desktop software absolutely independent from each other.

Examples of hypervisors are VMware ESX, Citrix XenServer and Hyper-V from Microsoft.

The concept of Virtual Desktop Infrastructure was influenced by three concepts: the traditional desktop, the traditional server-based computing using terminal servers, and server virtualization. The good thing with Virtual Desktop Infrastructure is that it gives you, in a manner of speaking, the best of all concepts that are the inspiration for it.

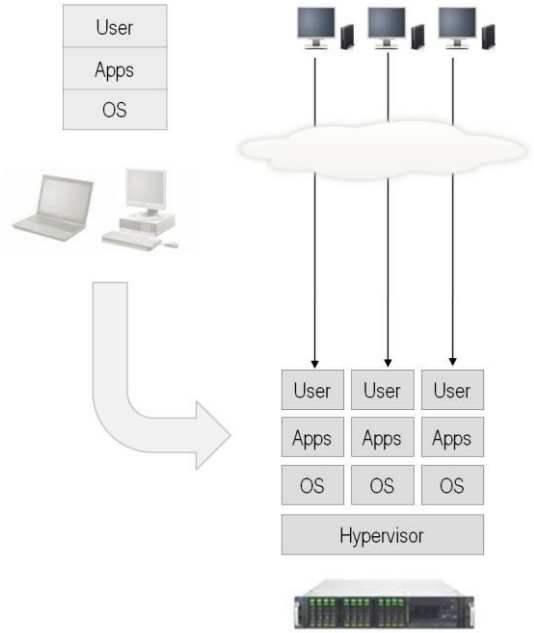

Customer benefits

Virtual Desktop Infrastructure yields a multitude of benefits for customers. First of all, desktop lifecycle management is tremendously simplified, because testing multiple desktop hardware configurations is no longer needed. Desktop deployment happens rapidly, as much as software updates and patches. Frequent desk-side visits are a thing of the past. And the retirement of virtual desktops does not cause more than a mouse click, without any data sanitizing at the endpoint.

Due to the absolute independency of hardware and software, hardware and software lifecycles are extended, thus enabling legacy software to run on latest hardware technology.

Moreover, business continuity can be maintained much better. If a virtual machine fails, it can be recovered very fast, which in turn reduces downtime. Even disaster recovery concepts can be realized, if there are end users with highly business-critical applications.

The fact that data is centrally safeguarded in the data center, reduces the risk of data theft. Likewise, the risk of data loss is reduced, because backup policies can be centrally controlled and enforced. As a result, virtual desktop infrastructures increase the level of security, and help fulfill compliance demands much better.

Due to the flexible access options, flexible working models are supported. And businesses have the flexibility to run their virtual desktops wherever they want; either on-premise in their own data center, or off-premise anywhere else.

Although additional server and storage infrastructure is needed in the data center, in an overall consideration, energy consumption and cooling requirements can be reduced, in particular, if you use thin or Zero clients as access devices. And don't forget that lifetime of thinner hardware is much longer compared to lifetime of a full blown PC which has a positive impact on costs, too.

The result of the benefits discussed is much better quality of service, high end user productivity and satisfaction, high flexibility and speed when you have to react on dynamic changes, and operational cost savings. All this in turn are the basic parameters for having a competitive advantage.

Connection brokering

Assuming you know the IP address of your virtual machine, you could manually enter the IP address in order to get connected. However, this would require that the virtual machine is always allocated to you and always started.

If the desktops of a larger number of users are hosted in the data center, it is questionable whether you need them all persistently. Some users are always sick, on vacation or on the move, or they don't need their virtual desktop for some other reason. Therefore, in order to save datacenter resources and money, it could make sense rather to create a pool of virtual desktops whose size is equal to the maximum number of simultaneously active users. Then the question comes up how a user, who wants to open a desktop session, will find his virtual desktop.

The dynamic connection of end users to their virtual desktops is the task of a connection broker. Users send their connection request to the connection broker either via an agent running in the access device or via a browser, just by specifying the connection broker's web address. After authentication, it is checked whether there is a dedicated virtual machine allocated to this user, or if the virtual machine from his last login is still available. If your personal environment is not available in a virtual machine yet, typically an available virtual machine will be selected from a pool of virtual machines and personalized according to the user profile. If there is no virtual machine available, the creation of a new virtual machine will be initiated, followed by the deployment of operating system and applications, as well as the personalization, e.g. using Microsoft's Active Directory. Then the connection broker will establish the connection. To improve security, the connection between access device and virtual desktop is encrypted.

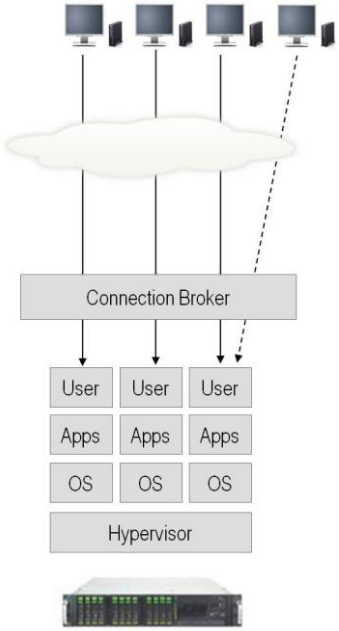

Connection brokers typically know all about available resources and take on the administration of virtual machines. They also look after reconnecting users to their desktops after disconnection. This is amongst others important for users who start a session at one thin client and want to continue it at another one next door.

Connection brokers typically offer the USB redirection feature, which supports the operation of local peripherals, such as storage devices, printers and scanners.

The most prominent connection broker products in the market are XenDesktop from Citrix and View Manager from VMware.
Microsoft's Remote Desktop Services (RDS) include a connection broker, too.

Typical use cases for VDI

Virtual Desktop Infrastructure addresses those scenarios, where an individual user always needs the highest flexibility with regard to using hardware and software resources, and where an interference with other users should be avoided. These are the typical characteristics of what analysts denote as knowledge workers.

After this generic specification, we are going to discuss some typical use cases or even compelling events, where the VDI concept helps.

Nowadays, mergers and acquisitions are a daily occurrence. As a result, new users have to be integrated with existing workplace environments. VDI tremendously accelerates this process and enables starting the integrated operation at an accurate point of time.

With an OS migration (e.g. from Windows XP to Windows 7) you always take the risk, that certain applications do not run on the new OS platform. Using application virtualization in combination with VDI basically reduces this risk down to zero.

Whenever a hardware refresh is planned, it might make sense to re-think your IT strategy. Going for VDI extends the hardware lifecycle and helps avoid huge investments in new workplace devices, while enjoying all the benefits shown before.

When introducing BYOD (Bring-Your-Own-Device), from an infrastructure perspective, VDI is an unconditional prerequisite.

A special challenge for IT organizations is the support of end users in remote and branch offices, or even on construction sites. Usually onsite support is very costly, and sometimes even almost impossible, notably when looking at factory plants in rural areas or foreign countries. Running such workplaces as virtual desktops in the data center removes all associated issues at one go.

Another typical example is software developers. While developing and testing new versions of applications, multiple restarts will happen every day. As this would be an encumbrance to other users incidentally working on the same server, it is advisable not to run a developer in a traditional server-based computing environment. However, more and more enterprises go for off-shore development with developers based in foreign countries. Using VDI, you can make sure that all development results, including all source code, never leave the data center.

Similarly, external staff, e.g. external project members, can easily be integrated and excluded from the customer's infrastructure at the end of the project.

In those cases, where data loss and downtime is business-critical, the high availability and disaster recovery capabilities of VDI are recommended. Enterprises which are strongly subject to legal restrictions and compliance demands can fulfill these much better when using VDI.

In training centers, training participants should have their own desktop for the duration of a training course. Using VDI, it is very simple to achieve this. After the course, all virtual desktops can be easily reset to an initial state by a mouse click.

With PC workplaces in dirty or dusty production environments, the failure rate, in particular of their rotating parts, is extremely high. VDI enables the usage of thin clients without rotating parts, and reduces maintenance cases and downtime dramatically.

VDI and traditional SBC in combination

If an existing terminal server infrastructure is supplemented by VDI, applications running on the terminal server farm are frequently needed by VDI users, too. In this case, you will have two options how to make these applications available for the VDI users. Applications can run locally in the virtual machines of these users, or the same applications running on the terminal server farm can be used. The latter option means lower resource requirements on the VDI side, leaner virtual desktops, more virtual desktops per server, and therefore lower infrastructure costs.

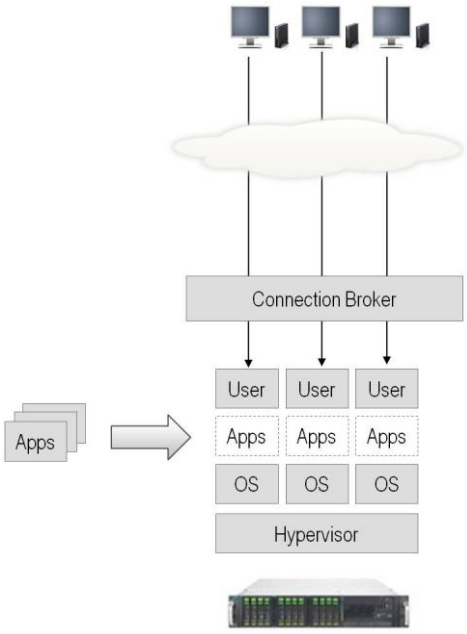

Desktop Virtualization is more than VDI

Using VDI you just separate the desktop operating system from the hardware underneath. You can increase flexibility and efficiency even more, if you separate the applications from the desktop operating system or the user profile from the desktop image. This brings technologies such as application virtualization and user profile virtualization into the game. Printing to any local printer requires the separation of the printer hardware from the desktop image, and requires therefore a printer virtualization solution. So, desktop virtualization is absolutely more than just VDI.

Moreover, to optimize desktop virtualization solutions, methods for storage reduction (e.g. OS streaming, volume cloning, data de-duplication and thin provisioning) need to be taken into consideration, just as network accelerators and components for multi-media and graphics support.

By a combination of these various technologies and solution approaches you may address all the questions that VDI leaves open. A more detailed insight into these technologies is given subsequently.

Application virtualization and application streaming

When installing applications, no matter if on a PC or server; no matter if on a physical or virtual system, entries will happen in the operating system's registry, and the file system is modified. Frequently asked exciting questions are:

- Do different applications interfere with each other?
- Will there be any DLL conflicts?
- Are registry entries, which originate from one application, overwritten by another one?
- Will other applications still run after uninstalling a certain application?

Finding a reliable answer often requires lengthy compatibility tests.

These efforts can be avoided by application virtualization. Instead of installing an application, a virtualized package with either a single application or several compatible applications is created and centrally stored. From there the applications can be streamed on demand to the target system, where they will run fully isolated in a sandbox with its own registry, and thus do not get in touch with the operating system. Therefore, even incompatible applications can be run in parallel on the same system. New software versions can be tested in parallel to the previous ones. Likewise applications can run on a system whose operating system actually does not support them. When the application is no longer needed, it can just be deleted and will not leave its mark. A de-installation is not needed.

Usually organizations take the effort of application compatibility testing for standard applications, which are used by a large number of users, while rarely needed applications are covered by application virtualization and application streaming, in order to avoid all the cumbersome and expensive quality assurance efforts. More and more, application streaming is used as a principal method of application delivery on demand. Organizations see a major advantage in the fact that applications are stored only once, and need to be updated and patched only once either.

Another application scenario could be mobile users who want to use one or a few applications while being on the move. As sandboxes can run everywhere, you just have to take your virtualized (sandboxed) applications with you, e.g. on a USB stick, and run them wherever you find a PC.

However, not all applications can be virtualized. These are, for example, applications requiring drivers or hardware dongles for security purposes. Except to this, it is important to know that, at least in some solutions, logically separated applications which communicate with each other in terms of a dynamic data exchange have to be virtualized in the same virtualization container.

Application virtualization and application streaming are not just good for virtual desktops, they are useful for all previously discussed popular paradigms. As every virtual layer eats resources, running virtualized applications in a sandbox causes a certain overhead and accordingly performance loss compared to installed applications. It is needless to say that for streaming a corresponding broad bandwidth is required. However this one might be available in the data center anyway.

Examples of application virtualization and streaming solutions are the Application Streaming component included in Citrix XenApp, Microsoft App-V (Application Virtualization), and VMware ThinApp.

OS streaming

Imagine a larger number of users with their virtual desktop environments hosted in the data center. As a start, this means that for each of these users a desktop operating system exists, which needs to be configured, stored, updated and patched. This will cause a large amount of work. If you then consider the amount of storage needed for all those samples of the operating system, you will soon come to the conclusion that storage becomes a major cost factor of virtual desktop environments. Therefore it is essential to have a look at methods that help reduce the required storage capacities.

The idea of OS streaming (sometimes denoted as image virtualization) is not to install the operating system on each system, but rather to make a single operating system virtually available in the network. For this purpose, you build one or few generic master images which are centrally deposited in the data center.

Virtual desktops are booted from one of these images over the network. After personalizing the virtual desktop according to the user profile, the user will be connected. Required applications which are not part of the image have to be deployed or streamed separately. The result of combining the technologies mentioned is a dynamic desktop delivery on demand.

OS streaming implicates an enormous potential of rationalization. Storage savings are enormous compared to storing dedicated images for each user. Let us consider a Windows 7 instance which is 20 GB in size. In an organization with 500 users this would mean 10 TB of storage in total. In a large enterprise with 10,000 users in total 200 TB would be required, just for the virtual desktops. Using operating systems streaming from one central operating systems image, only 20 GB are needed, Assuming that 5 different images for 5 different user types exist, let it be 100 GB, a huge difference compared to 200 TB.

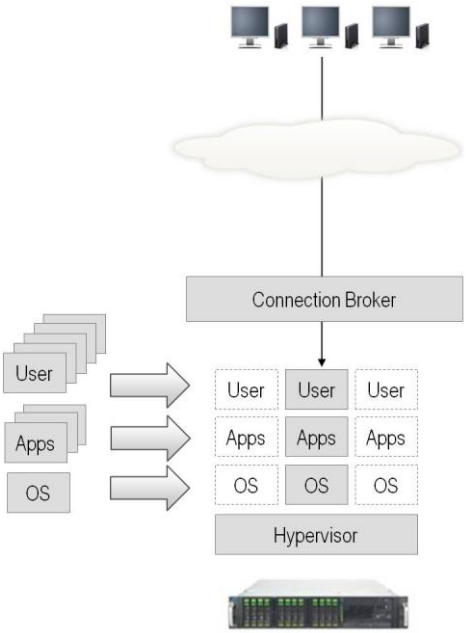

With OS streaming, basically all systems are stateless and easy to manage, because they don't need any updates and patches. There is only one or few operating systems instances that need to be created, tested, updated, fixed and patched. Scaling is extremely easy. Infrastructures can be flexibly extended just by adding hardware, that's it.

The ideal case is of course to have only one master image. However, this requires that all users will always use the same version of the operating system and the applications.

OS streaming is not limited to virtual desktop infrastructures, of course. Similar to application streaming, OS streaming can be applied to traditional PC workplaces and traditional server-based computing as well.

An example of an OS streaming solution is Citrix Provisioning Services.

Volume cloning

Another approach to reduce storage space is volume cloning. Similar to OS streaming, in a first step, a master image is created. Deploying virtual machines happens very rapidly by creating so-called linked clones, which all share the master image.

Initially the linked clones are empty virtual disks, which except to a reference to the golden master, do not require any storage space. Linked clones grow dynamically when user-specific changes occur to individual virtual machines, e.g. by installing new applications or creating files on the c-drive.

For patches and updates, only the golden master image needs to be touched. Once this happens, the linked clones will be deleted, newly created and simultaneously rolled out. All changes to the individual desktops will then be lost. For this reason, it is strongly recommended that user-specific data is stored on another volume or a private network share.

An example for volume cloning is the FlexClone technology from NetApp, which is implemented on the storage level using storage-based snapshot technologies for acceleration. In contrast, View Composer from VMware, which is deeply integrated with View Manager from VMware, is a software solution, absolutely independent from the storage system being used. Another example is Machine Creation Services from Citrix.

Data de-duplication

OS streaming and volume cloning are focused on reducing system storage space. Especially when it comes to reducing disk storage space for user-specific data, data de-duplication for primary storage becomes interesting. With this method, data duplicates are identified and discarded. For applications and users, data de-duplication is fully transparent.

Of course, data de-duplication is not exclusively bound to user-specific data; it can be applied to system storage, too. However, be aware that without OS streaming or volume cloning, you still need to manage a multitude of virtual desktops.

Even a combination of volume cloning and data de-duplication can be attractive. Due to volume cloning, only one or a few images need to be managed and the initial storage capacity needed for the cloned virtual desktops is quasi equal to zero. When clones are increasing in size, you should assume redundancies in various clones as well, which can be eliminated by data de-duplication, in order to keep the overall storage capacity requirements on a low level.

Data de-duplication for primary storage is a function implemented for instance in disk storage systems from NetApp. It should be mentioned that backup on the file level would actually mean a multiple transmission of data duplicates over the network and a multiple depositing of these data on secondary storage.

Admittedly there are alternative de-duplication technologies which de-duplicate data at the source before a backup. This causes data duplicates to be transmitted over the network and deposited on the secondary storage only once. An example is Avamar from EMC.

Thin provisioning

Whenever you created a new volume or data partition in the past, the maximum amount of physical storage, which would be needed during its entire lifecycle, was immediately allocated to a user, application or system. The result of this procedure, which was called hard provisioning or fat provisioning was high storage capacity requirements and poor utilization. In various studies people speak of 25% utilization at average, especially when it comes to user data volumes.

The so-called "Thin provisioning" technology can change the situation. Thin provisioning separates the physical allocation of storage space from the virtual storage space which is visible to the user. Only when writing new data, free storage space out of a storage pool is allocated. Nevertheless, maximum capacity is always guaranteed. Just in case the allocated overall storage capacity exceeds a pre-defined threshold, the administrator will get a notification that the physical storage pool will have to be expanded.

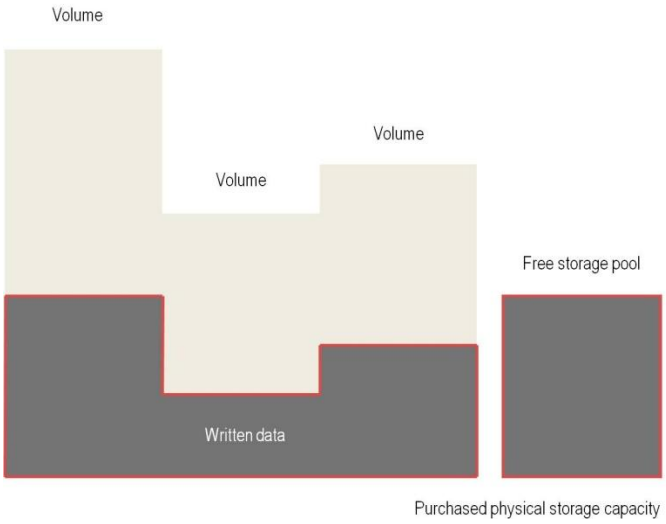

Thin provisioning reduces storage capacity requirements tremendously, which in turn impacts on floor space and energy consumption, management efforts and maintenance, and of course total cost of ownership. People speak of a 50% reduction of infrastructure cost and a 90% decrease of personal-related costs. In addition, storage utilization is enormously improved.

Thin Provisioning comes integrated with storage products from Fujitsu and NetApp. VMware implemented thin provisioning fully software-based in their View solution enabling the same benefits also with storage systems shipped without the thin provisioning feature.

User virtualization

Perhaps one of the most challenging aspects of the PC workplace from the IT management perspective is the unique nature of the PC. As soon as employees receive their PC, they will transform that standard desktop into one which is unique to them. The more personal the traditional PC workplace, the more difficult it is to manage, since personal changes at the operating system and application level move the desktop out of its original standard and predictable form.

Desktop component model

This is one of the reasons why organizations are looking at the component model of the desktop, essentially separating the things that make up a PC workplace into standardized, predictable assets, such as corporate OS and applications, and unique, user-centric assets, such as policy and personalization settings.

By managing all aspects of the user independent from the desktop, IT organizations are able to standardize the corporate operating system and applications, delivering them on-demand only when needed. This method enables companies to eliminate unnecessary desktop management costs while ensuring users of all types receive the very best working experience - even in the most heterogeneous environments.

By automating the delivery of standard operating system and applications from a single image and then dynamically applying the user environment, thousands of personal workplaces across the enterprise can be created on a low-cost level.

The user environment typically consists of two elements, managed together from a user-centric perspective. These are the personalization of the user and the enforcement of user policies.

Personalization

Personalization constitutes any change a user makes to his desktop. In the past, user changes were stored in the user's profile, which is a group of settings that define the environment to be loaded when a user logs on. With a traditional PC, the user profile is stored on the local hard disk. Any changes to the local user profile are then specific to the PC on which the changes are made.

In a virtual desktop infrastructure, a local profile will not make too much sense. Especially when users are connected to different virtual desktops of a virtual desktop pool, a roaming user profile is more suited. The roaming profile is stored on a network drive and downloaded to a virtual desktop, every time that a user logs on. Any changes made to a roaming user profile are synchronized with the central copy when the user logs off.

Roaming user profiles can bloat to 100s of MB in size. Downloading them at logon can cause heavy network traffic and performance degradation. It is true that a profile can be cached on a virtual desktop, so that the user could have a good experience when logging on the next time. However, dynamically connecting a user to a virtual desktop on demand essentially gives a new virtual desktop every time the user logs on. As there is no cached profile then all of the users profile will have to be downloaded again. This issue is only extrapolated when adding lots of users logging on at the same time.

Furthermore, roaming profiles have a high requirement for support. For instance when the network connection fails during a profile save or when a user roams between multiple desktops with slightly different software versions, application settings can easily become corrupt.

The high network load could be circumvented by a mandatory user profile which contains consistent settings for all users and can therefore be local to each virtual desktop. It is the administrator who predefines the mandatory profile, and is the only one who can make changes. Any changes made by the user to desktop settings are lost when the user logs off. This is to say, that individuality is not supported at all, something that is unsuitable for the majority of user cases for virtual workplace solutions.

Virtual user profiles combine the performance advantages of mandatory profiles with the flexibility offered by roaming profiles. They assure that only relevant and changeable personal settings are copied over the network at logon and logoff. This reduces logon times dramatically. And the potential for profile corruption is removed, too.

An example of a virtual user profile solution is Citrix ProfileManager. In particular, AppSense Environment Manager should be emphasized as a personalization solution which is part of their user virtualization offering, because in that solution personalization is abstracted from a desktop, stored separately from a user profile and as such is now able to be applied to any image on demand, and not all at logon. The impact is that network load is further reduced along with enabling personalization settings to roam between different desktop types.

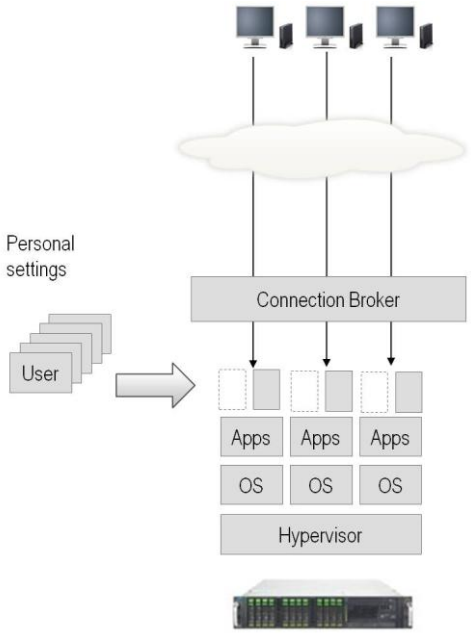

User policy and enforcement

User policy is used to set up and maintain a user desktop session. Policy also ensures a user session remains compliant by controlling application access, locking down or removing operating system and application functions, even self healing essential files, folders, processes, services and registry settings.

For this purpose people frequently use logon scripts which are executed, when the user logs on. Subsequently the policies remain static and are valid during the user's entire session. With a traditional PC that is always at the same location this might be adequate in many cases. In the event of virtual desktops in the data center which can be flexibly accessed from anywhere and any device, the question is if the policy should always remain the same, if for example the user moves to another location during his session?

Looking at common scenarios like a financial trader or a medical worker as examples, should they always have access to the same application, printer and network drives, or should this be dynamically provided during the session based on certain criteria or conditions? When the trader moves away from the trading floor, it should be possible to remove the access to his trading applications, in particular when there is a legal requirement. Similarly, if the medical worker is connecting in from outside of the hospital, perhaps we would want to remove the ability for him to access certain applications, medical data or remove the ability to print confidential information.

Using a "logon only" method for assigning policy, the user would still be able to access these functionalities, if the user logged on in one location and then moved to another location within the same session.

To avoid these problems, solutions are needed, with which you can adapt user policies to new situations dynamically during a session. This can be achieved by using trigger points, conditions and actions to define the policy. Besides logon / logoff and start-up / shutdown of the desktop, trigger points can be used to apply policy at process start / stop, network disconnect / reconnect and session locked / unlocked. Trigger points can be seen as the when something gets done, however only based on certain conditions. Conditions are related to the questions, e.g. who, where from or how a user is connecting to a virtual desktop or application. Actions resulting from these triggers and conditions include file, folder, registry, ADM files, drive mappings and printer mappings.

Let us regard a few examples. Assuming you move from one floor to another floor in your office building. At the moment you are reconnecting to your virtual desktop, the nearest printer to your new location could be mapped as the policy could automatically be re-evaluated when the user reconnects to the session. Likewise you can achieve, that certain network drives or printers will be mapped only when an application is started for which these resources are relevant. The flexibility of policy enforcement is only limited by the imagination of the administrator.

By easily manipulating triggers, conditions and actions, an administrator can easily and quickly implement business policies which can be shared and utilized across operating system boundaries and different application delivery mechanisms.

AppSense Environment Manager is a complete user virtualization solution including both policy enforcement and personalization. Policy and personalization are abstracted away from the desktop and managed independently from the operating system and the application layers. The separate layer is applied on demand to configure and personalize the image regardless of how the operating systems or applications are delivered.

Printer virtualization

On a traditional PC, especially on a mobile device, various printer drivers need to be available in order to enable printing at various locations, be it in the office, at home, or anywhere on the move. As virtual desktops can be flexibly accessed from anywhere, the user's desire is to print anywhere as well. But as you can never know which printer will be next to any access device, it will be difficult to be always equipped with the right printer drivers.

Printer virtualization breaks the tight coupling of the desktop image including operating system and applications from printer hardware. No specific printer driver is needed any longer in the desktop image; this makes driver-free printing reality.

Moreover, printer virtualization solutions usually address the problem, that in any server-based computing scenario, print jobs related to local printers could block the wide-area networks, which in turn can have a negative impact on interactive users. Due to highest data compression, reliable printing at a high speed is ensured, even when bandwidths are limited. Data encryption makes printing extremely secure.

An example of a printer virtualization solution is .print from ThinPrint.

Network acceleration

In a virtual desktop infrastructure, there is always a network between the virtual desktop and the access device. The required network bandwidth is strongly influenced by various parameters. These are the applications and the content to be transmitted, the resolution and the number of displays at the access device, local peripherals, such as mass storage or printers, the end user behavior, and the number of users per connection. Especially in wide area networks with low bandwidths or long distances, it is imaginable that virtual desktop infrastructures combined with the technologies mentioned earlier, initially do not fully meet all performance expectations. In these cases further technologies for optimization should be considered.

In order to minimize network bandwidth in general, there are WAN accelerators which are aimed at reducing and compressing data to be transmitted over the network and at protocol optimization.

Lack of network bandwidth can also be overcome by cached storage, where virtual desktops of users in a subsidiary are replicated in the subsidiary and synchronized in regular intervals with the data center. It is just the virtual desktop infrastructure in the data center which will be managed.

Of course, user profile virtualization and printer virtualization solutions, as discussed earlier, contribute to reducing network load, too.

Graphics and multi-media applications

A particular challenge with virtual desktop infrastructures is graphics-intensive applications and multi-media applications, such as video and audio streaming, or real-time communication. While in the past, these application types were not appropriate for server-based computing, meanwhile solution approaches are in place, which are sufficient for a broad diversity of scenarios.

In the event of graphics-intensive applications, the bitmap generation, i.e. the rendering, can happen on either the terminal or the server side. With terminal-side rendering, graphics instructions are sent to the access device, where they are interpreted and executed either by a local graphics chip or by software running on the local CPU. With server-side rendering, the complete digital image is built up on the server and then sent to the terminal to be displayed. In both cases, data will be compressed by the remote desktop protocol before transmission and decompressed thereafter.

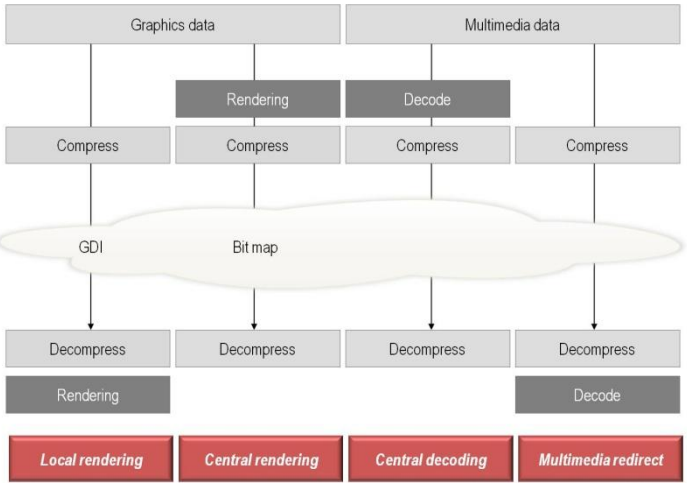

In the event of multi-media, content can also be decoded centrally by running a media player in the virtual desktop. The remote desktop protocol will look after the compression of the decoded content and its decompression locally at the access device. As decoding multi-media streams can heavily load server resources, multi-media redirection can be used to move compute load from the server to the access device. In this case, encoded multi-media content is compressed by the remote desktop protocol and transmitted through the network. At the access device, it will be decompressed and then decoded. Doing so, multi-media content reaches access devices efficiently at a high speed. Multi-media re-direction should only be enabled, if the access device has sufficient resources to handle local multimedia decoding. Otherwise it should be disabled by the administrator

Multi-media applications running in virtual desktop environments often pose the question about user experience and user satisfaction. This question cannot be answered across-the-board, because here content complexity and content quality matter. Let us take video as an example. As in consecutive frames only delta information is transmitted, indoors pictures will cause only low volume data streams. In contrast, complex and moving pictures from outdoors, e.g. from a soccer match, will cause extremely high volume data streams. Content quality is influenced by the display resolution, the frame rate and the compression. And user satisfaction depends on the expected quality, which in turn influence the price of the overall solution.

IP telephony

Nowadays, more and more customers transact telephony via the Internet and use voice-over-IP software solutions running on the PC systems of their employees, enabling phone calls by means of a headset connected to the PC.

IP telephony means a tremendous cost reduction, because all phone calls are basically free of charge. Furthermore such solutions are frequently used in conjunction with integrated telephony applications, which on the one hand are to identify the calling person by means of the phone number, and on the other hand are to support dialing by a mouse click.

It is obvious that these customers also want to enjoy the benefits of IP telephony in a virtual desktop infrastructure. For this purpose, the access devices need to be equipped with interfaces for the headset (either USB or audio line-in and out) and a sound card, while the connection broker needs to support bi-directional audio. The voice-over-IP software will typically run in the respective virtual desktops.

Hosted Central Desktop

In those cases, where IT workplaces with graphics applications are to be centralized, but a virtual desktop infrastructure will not provide the desired user experience, graphics workstations can be moved into the data center and accessed remotely from a thin access device. This concept is known as “Hosted Central Desktop”.

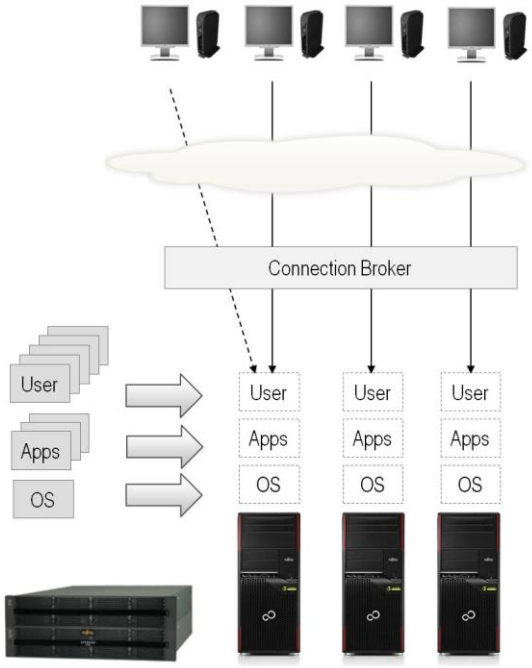

What are the reasons for centralizing graphics workstations? The main reason is definitely security and the strong demand to protect intellectual property. In the event of a maintenance case, support staff requires access to the workstation. If the workstation is in the developer’s office, they could easily catch confidential information (e.g. details of a development) and pass this to a competitor. Having the workstation in the data center, you can effectively avoid this problem, because there is no more need for support staff to enter the developer’s office. Moreover, confidential information is centrally safeguarded, significantly reducing the risk that any unauthorized data copy is generated. If it is required to present development information necessitating a workstation at a different location, this can be achieved without having to physically transport a workstation, avoiding risks of theft, breakage or confiscation in the case of crossing country borders.

Another reason for centralization is to avoid the heat and noise generated at the user’s desk, which can compromise productivity. Having workstations located in an office frequently requires a powerful and costly air conditioning. Besides, centralization creates more space in the developer’s office, and enables a flexible access to the workstation from anywhere.

Pooling of graphics resources

Moving workstations from the user desks into the data center opens up other opportunities, too. Frequently, developers do not require all the expensive graphics resources all the time. For resource optimization purposes, a pool of workstations can be established in the data center, the number of workstations being lower than the total number of users. Using a connection broker at logon, an available workstation would be selected and allocated to the user. If all physical workstations are being used, a new user wanting to logon will not be able to, as a workstation cannot be allocated. This is basically the only difference from a virtual desktop pool.

The benefits are obvious. Less workstations means less acquisition costs, less maintenance costs, less energy and less operating costs.

Physical workstation pooling requires the separation of the user personality from the physical workstation. In other words, a roaming profile must be used and user data must be on a central storage system.

Local Virtual Desktop

All solution approaches discussed by now require a connection from the end user's device to the virtual desktop environment in the data center. As long as server-based computing has existed, an essential issue has been mobile users. It is true that there are more and more technologies available for getting connected while being on the move. The number of WLAN access points is steadily increasing, and so does the bandwidth of wireless wide area networks.

However there are situations where all these networks are not accessible. Typical examples are when you are traveling in a train where usually even mobile phones don't work, or when you are traveling in a plane without a WLAN.

If you want to use travel time in the train or plane to do some work, it might be advantageous, if you could take your virtual desktop environment with you. This is exactly the idea of local virtual desktops, which is sometimes also denotes as "Offline VDI", which allows mobile users to run multiple virtual desktops on their notebook. As a prerequisite, a hypervisor is required on the notebook, either running on the host OS of the notebook (type-2 hypervisor) or on bare metal (type-1 hypervisor). Examples of type-2 hypervisors are VMware Player and VMware Workstation which run on every Windows and Linux OS. Another example is Microsoft MED-V. Examples of type-1 bare-metal client hypervisor is XenClient from Citrix and Client Hyper-V from Microsoft.

Combining local virtual desktops with VDI, the virtual desktop is once delivered from a central image to the notebook. Of course, all work done offline will only have an impact on your local copy. As soon as you get connected to the corporate network, your updates will automatically be synchronized with your virtual desktop environment in the data center, as system updates and patches will affect your local virtual desktop. The synchronization eliminates the need to backup notebooks, and the automatic update ensures that users always work with the latest software versions and security patches.

Virtual desktops are encrypted and fully isolated from each other or the host environment, if there is one. Additional security is provided by allowing policies to be put in place. For example, if a laptop hasn't re-connected to the corporate network for a certain period of time, the image will lock itself down. Likewise, a kill pill can be issued for stolen notebooks, which disable the device when it gets connected the next time. Thus it is guaranteed that the company's security policies are not compromised.

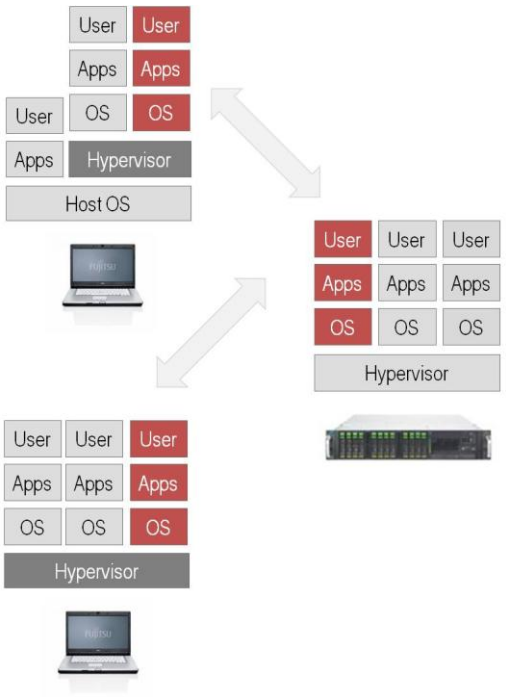

With local client virtualization as a supplement of VDI, IT organizations have workplaces of mobile users better under control than ever before, even without knowing how the end user's device looks like. And what is even more important: They are able to manage IT workplaces of mobile users in exactly the same way as they manage the virtual workplaces of their stationary users.

And what is in for the mobile user? Everybody is aware of the terrible impact when a notebook crashes or is stolen while you are on the move. With VDI and local client virtualization in combination, you may flexibly access your virtual desktop environment in the data center immediately, anywhere from any access device. And there is always the option to procure a new notebook and download your virtual desktop during the subsequent night, if you have access to your corporate network with a reasonable bandwidth. And the next morning, you will be able to proceed in exactly the same manner as before the crash or theft, without any complex system and application installation or data migration, and without any need to involve a support engineer.

Of course, you can even carry your entire virtual desktop with you on a portable or removable storage device, e.g. a USB stick with sufficient capacity. And on the move, you can run it on any PC with the respecting hypervisor support. It is even imaginable to have the hypervisor on the storage medium.

Type-1 or type-2 hypervisor?

Using a type-1 hypervisor, basically the full performance of the host device is available for the virtual desktop environment. No other operating system sits in between which would eat additional resources. Due to the fact that virtual desktops are fully isolated from each other, you will achieve the highest level of security. However, existing devices can only be used, if the hypervisor is released for them.

With a type-2 hypervisor which is usually available for the most common desktop operating systems, existing devices can simply be used to run virtual desktops. No special device drivers are needed, because the virtual desktops will use the device drivers of the host operating system. But it is a matter of fact, that the host operating systems needs system resources, meaning that the individual virtual desktop will by far not benefit from the performance the device as such makes available.

Use cases for local virtual desktops

Besides any type of mobile user, local virtual desktops are applied whenever multiple roles of an employee require multiple workplace environments. For example, software developers need a development and test environment in addition to their normal end user environment; likewise, a system administrator might use different environments depending on whether he acts as an administrator or as a normal end user. Here it is worth mentioning that local virtual desktops can also run on stationary client devices, such as desktops or workstations.

A new trend: BYOD (Bring your own device)

The combination of local client virtualization and VDI enables mobile users, for the first time ever, to use their own private notebook inside the company without violating any security policies. You may run your private environment and your business environment as different virtual machines on the same device. In the US already today many companies give their employees a certain amount of money per year for a private notebook which is also used as a vehicle to run the virtual desktop environment made available by the organization. The IT department will only manage that virtual desktop; anything else is up to the end user.

Local Streamed Applications

An alternative for offline usage is Local Streamed Applications. Business applications are once downloaded to the mobile device where they will run in a sandbox. Data used or generated by the applications can be totally isolated and separated from what else is on the device. For the rest, all security mechanisms known from the Local Virtual Desktop, such as data encryption, enforcement of policies and remote wipe are also available.

Local Streamed Desktop

A virtual desktop infrastructure in the data center simplifies management, but requires a certain investment and implementation effort. By applying the previously discussed concept of OS streaming to physical PC workplaces, you can simplify management without building such a virtual desktop infrastructure in the data center. This concept is also known as "Local Streamed Desktop".

The idea is to use a diskless PC as working device, stateless and without any operating system installed. Instead of that, the operating systems image resides on a central storage, and will be booted through the network, whenever you want to start your session. All applications will be executed locally.

The local execution of all applications leads to an optimum user experience, and local peripherals can be used as usual. Local streamed desktops can be introduced rapidly, and will lead to simplified management at low infrastructure requirements. Especially if a standard image is used for many users, storage capacities are tremendously reduced.

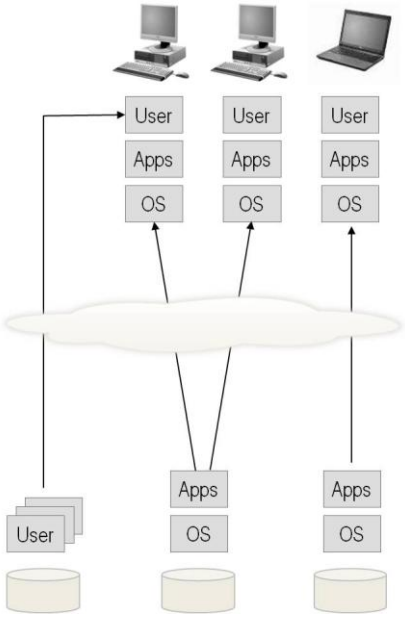

However, streaming the entire image through the network requires an enormous bandwidth between your PC workplaces and the data center. That is why local streamed desktops are not appropriate, if a WAN separates the office from the data center. As you can use any local peripherals, it is quite easy to export data, thus being faced to the same security risks you have got with a traditional PC.

Web Desktop

In the last couple of years the Web has become the main workspace for many users. More and more of the applications needed to do their work are web-based, or at least accessible through the web.

The Web Desktop becomes the aggregator for these applications. For accessing web-based applications, an HTML5 compatible browser is sufficient, which will be available on any device, no matter which operating system is deployed.

Which model for which type of user?

After having discussed various virtualization models, we should turn towards the question which model to use for which type of user. There is certainly no universal answer, but indications which might be applicable in many cases.

For task workers using only a few standard applications, the Hosted Shared Desktop is sufficient as a rule. For knowledge workers with higher demands regarding personalized environments, the Hosted Virtual Desktop is recommended. Power users with high performance requirements should have their own PCs or workstations. If centralization is demanded, the Hosted Central Desktop would be the optimum solution. For all these models it is recommended to use thinner devices for the access. Mobile workers who are frequently on the move should be able to use their workplace environment offline. Therefore a mobile device with local virtual desktops or Local Streamed Applications could be the solution. Choosing the right model for external users mainly depends on the specific tasks that need to be done.

Application and user profile virtualization are always beneficial, no matter which of the other models is applied.

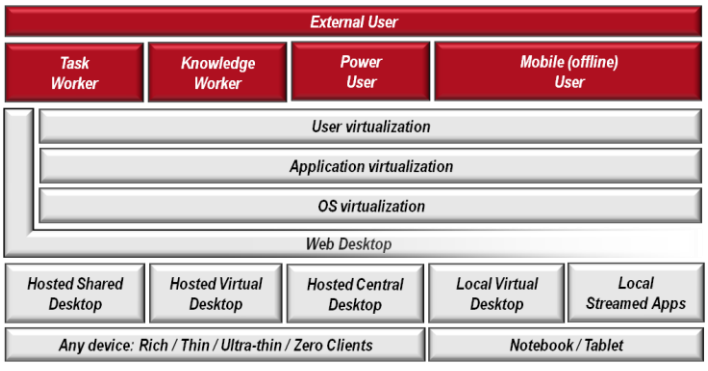

And for those users, who mainly need access to web-based applications, the web desktop will be an efficient solution, if network coverage can generally be assumed.

An important finding is the fact that one size does not fit all. Since user diversity is quite high, enterprises need an intelligent combination of diverse models. What is more, economy does matter. And it should not seem that any of these virtualization technologies is a must in all cases. By implication, it can turn out that for certain scenarios the traditional approach with just using a stationary or mobile device can be economic as well.

Fujitsu's approach – One stop shop for Virtual Client Computing

In the field of virtual client computing, an ever-increasing number of technology choices can be considered in crafting a solution that optimally meets the specific requirements for the organization. For most customers, certainly not all of them are relevant. However, understanding which technologies bring the most beneficial results is an increasingly complex and time-consuming task.

Selecting the appropriate technologies according to the customer's business requirements, evaluating, testing and combining all these technologies, as well as finding the right mix is a new challenge for customers, which should not be underestimated. The fact, that almost all technologies are rather new, does not make the task easier, as the skills to evaluate all these technologies need to be built first, and this can take considerable time and effort. Moreover, the various technologies originate from an ever-increasing pool of vendors, which in turn increases the complexity even more.

Similarly, even after the optimal recipe has been decided upon, the integration of numerous building blocks, such as servers, storage systems and access devices, virtualization middleware, desktop operating systems, management software for virtualization, sometimes even traditional workplace management tools, and applications means that the work has just begun. Uncertain project duration and a multitude of risks can be the consequence.

This is exactly where Fujitsu comes into the game. Fujitsu's virtual client computing approach is concerned with taking out the complexity, reducing the customer's work effort, and reducing the overall risk in putting all the pieces together, with the goal of making IT simple for our customers, and helping them overcome the many hurdles in realizing a successful solution.

Market-leading virtualization middleware

Close partnerships with all prominent market leaders enable us to use best in class virtualization middleware and additional technologies to optimize the overall solution. Fujitsu can provide the respective licenses, the subscription advantage and the support.

Proven infrastructure products

Fujitsu's infrastructure products, such as PRIMERGY servers, ETERNUS storage systems, network components, FUTRO thin clients, the ultra-thin clients for remote workstation access, as well as the differentiating Zero Clients and Portable Zero Clients, represent an excellent basis for this purpose. These products are certified for all market-leading virtualization products, and have proven success in innumerable virtualization projects. The same is true for storage systems from our storage partner NetApp.

Optimized infrastructure bundles

For certain scenarios there are even infrastructure bundles available, which enable a virtual client computing solution out of the box. Optimized hardware configurations, resulting from Fujitsu's vast project experience, are pre-installed with the required middleware and software, and therefore ready to run.

End-to-end services – Consulting, design, implementation and support

Through all our activities, no matter whether in our labs or in real-life projects, we have gained experience as to what is required to successfully introduce virtual client computing solutions. This broad knowledge of optimizing solutions for specific customer requirements is reflected in our end-to-end services.

The realization of the optimum solution is typically not a trivial task, as organizations have unique business and IT structures and processes, which in turn reflect unique requirements for an optimum virtual desktop solution. Important parameters influencing the solution are the customer's business demands, the existing infrastructure, the customer's strategies, the overall costs and more. As Fujitsu works with all leading virtualization technology partners, it supports its customers with a neutral consulting approach, acting as a valuable advisor.

Using a structured process, based on extensive experience in real customer projects, concrete business goals are mapped to technology and solution criteria, to ensure that concepts are proved end-to-end. This reduces risk and accelerates informed decision-making.

In realizing customer-specific proof-of-concepts, Fujitsu partners with customers and offers them valuable consulting experience. Fujitsu's support encompasses guidance, coaching and participation in solution implementation and migration.

Optional complementary ROI services for IT investment decision support assist in formally determining the financial impact of infrastructure changes already in early project phases. Based on the maturity level of a project, the financial impact is either estimated or can optionally be predicted with a binding agreement. The ROI services are based on trademarked and proven methodologies, which have been successfully used in a wide variety of customer IT infrastructure engagements.

Naturally, Fujitsu's end-to-end services cover, in addition to consulting and design of the future architecture, the implementation, integration and maintenance of the overall infrastructure solution. And if financing options are looked for, Fujitsu will be the right partner as well. In other words: Fujitsu can provide virtual client computing from a single source.

All these elements contribute to reduce complexity, project time and risk, thus providing the foundation for a smooth implementation and operation, to enable customers to realize the value of a solution as quickly, and effortlessly as possible in support of their business and IT goals.

"Managed Workplace" services and "Workplace as a Service"

Besides implementing virtual client computing solutions, Fujitsu offers even alternative sourcing models for virtual desktop infrastructures. Fujitsu's managed service approach is targeted at customers who intend to rather concentrate on their core business rather than annoying routine tasks. These managed services represent a combination of our extensive experience in the lifecycle management of physical client devices and our managed data center services. Customers take advantage from a high level of standardization and a well-managed workplace infrastructure, scale effects which Fujitsu can offer as a service provider, the simple opportunity to alleviate shortages in resources and skills, flexible customer-specific and business-related service levels, and cost reductions. The "price-per-seat" charging model eliminates investment risks and ensures highest cost transparency. At the same time, customers keep their IT infrastructure fully under control.

Alternatively, with "Workplace as a Service", standardized IT workplaces can be delivered as a service from the cloud with a standardized service level agreement. The required desktop infrastructure is operated by Fujitsu on its own responsibility. A "pay-as-you-use" model is the basis for billing, minimizing capital expenditure.

Moreover hybrid models are imaginable, in which customers operate their workplace infrastructure themselves, but want to source certain applications as a service from the trusted Fujitsu cloud.

Summary

Virtual Client Computing helps customers improve service quality and reduce costs. However, selecting, evaluating, implementing and using the right technologies for an individual situation can be a complex task.

Fujitsu helps overcome the many hurdles in realizing a successful solution. In a structured process, starting with neutral consulting, Fujitsu designs an optimum infrastructure solution for the customer, based on best-in-class virtualization technologies and proven infrastructure products, and supports its customer in proving the virtual desktop goals established beforehand. Furthermore customers benefit from Fujitsu's experience in realizing virtual desktop infrastructures and a rapid implementation.

"Managed Workplace" services enable a careless operation, and using "Workplace as a Service" you will get your highly standardized desktop as easily as electricity from the socket or water from the tap.

All this helps reduce project time and risk. In a nutshell, in the same way as with other building blocks of its portfolio strategy, Fujitsu will help customers improve flexibility, efficiency and quality.

Contact
FUJITSU Technology Solutions GmbH
Address: Mies-van-der-Rohe-Strasse 8, 80807 Munich, Germany
Phone: +49-7203-922078
Fax : +49-821-804-88429
E-mail: gernot.fels@ts.fujitsu.com
Website: www.fujitsu.com/fts