Terminal server

From Wikipedia, the free encyclopedia

A **terminal server** is a specialized computer which aggregates multiple communication channels together. Because these channels are bidirectional, two models emerge: Multiple entities connecting to a single resource, and a single entity connecting to multiple resources. Both of these models are widely used. Both physical and virtual resources can be provided through a terminal server: centralized computing can provide multiple users access to a remote virtual operating system. Access Providers often use terminal servers to terminate physical connections to their customers, for example customers who get Internet through some form of modem.

Contents

[hide]

History

Historically, a terminal server was a device that attaches to serial RS-232 devices, such as "green screen" text terminals or serial printers, and transports this traffic via TCP/IP TELNET, SSH or other vendor-specific network protocol (e.g. LAT) via an Ethernet connection.

Digital Equipment Corporation's DECserver 100 (1985), 200 (1986) and 300 (1991) are early examples of this technology. (An earlier version of this product, known as the DECSA Terminal Server was actually a test-bed or proof-of-concept for using the proprietary LAT protocol in commercial production networks.) With the introduction of inexpensive Flash memory components, Digital's later DECserver 700 (1991) and 900 (1995) no longer shared with their earlier units the need to download their software from a 'load host' (usually a Digital VAX or Alpha) using Digital's proprietary MOP protocol. In fact, these later terminal server products now also included much larger Flash memory and full support for the TELNET part of the TCP/IP protocol suite.

Many other companies, such as Spider Systems, entered the terminal server market with terminal servers pre-loaded with software fully compatible with LAT and TELNET. Some manufacturers also stated specifically that they had emulated Digital's command set for terminal server management. Besides retaining the ability of the older terminal servers to obtain their run-time code from a load host, most were able to bootstrap from on-board flash memory or from a floppy disc held in a drive in the terminal server. Some Xyplex terminal servers could act as load host for each other; one would hold the code on a PCMCIA Flash card and serve it to another.

Starting in the mid-1990s, several manufacturers such as U.S. Robotics produced "modem terminal servers". Instead of having RS-232 ports, these would directly incorporate an analog modem. These devices were commonly used by Internet service providers to allow consumer dial-up. Modern versions interface to an ISDN PRI instead of having analog modem ports.

Modern usage

The term "terminal server" is used in three main ways, as of 2006:

- Console servers (also known as *serial terminal servers*) are often used for connection to the console ports of Unix servers. This then allows system administrators to connect to the servers over the network. This is important for rebooting the system and for hardware debugging, where the operating system will not boot correctly.
- A terminal server may refer to a network access server.
- Most commonly, "terminal server" means a server used in centralized computing (see next section).

Centralized computing

There are two contemporary models of centralized computing:

In one, the terminal server provides a Windows or Linux desktop to multiple user terminals - the modern term for these terminals is "thin client".

In the other model, an ordinary computer acts *temporarily* as a terminal server, providing its desktop to a remote computer over a wide area network such as the Internet, in order to enable teleworking.

A client of a terminal server is referred to as a thin client. Software clients used in the teleworking model are known as remote desktop applications; however, these remote desktop applications are also used in the thin client model as well.

> *For more details on this topic, see thin client.*

Thin client

From Wikipedia, the free encyclopedia

Jump to: navigation, search

This article needs additional citations for verification.
Please help improve this article by adding reliable references. Unsourced material may be challenged and removed.
(July 2008)

A HP T5700 thin client, with flash memory

A Neoware m100 thin client.

A **thin client** (sometimes also called a **lean** or **slim client**) is a client computer or client software in client-server architecture networks which depends primarily on the central server for processing activities, and mainly focuses on conveying input and output between the user and the remote server. In contrast, a thick or fat client does as much processing as possible and passes only data for communications and storage to the server.

Many thin client devices run only web browsers or remote desktop software, meaning that all significant processing occurs on the server. However, recent devices marketed as thin clients can run complete operating systems such as Debian Linux, qualifying them as diskless nodes or hybrid clients. Some thin clients are also called "access terminals."

As a consequence, the term "thin client", in terms of hardware, has come to encompass any device *marketed as*, or *used as*, a thin client in the original definition – even if its actual capabilities are much greater. The term is also sometimes used in an even broader sense which includes diskless nodes.[1]

Contents

[hide]

[edit] Introduction

The thin client is a PC with less of everything. In designing a computer system, there are decisions to be made about processing, storage, software and user-interface. With the reality of reliable high-speed networking, it is possible to change the location of any of these with respect to the others. A gigabit/s network is faster than a PCI bus and many hard drives, so each function can be in a different location. Choices will be made depending on the total cost, cost of operation, reliability, performance and usability of the system. The thin client is closely connected to the user interface.

In a thin client/server system, is placed the software for the user interface on the thin client, possibly some frequently/heavily used application, and a networked operating system and nothing else. This software can be loaded from a local drive, the server at boot or as needed. By simplifying the load on the thin client, a thin client can be a very small, low-powered device giving lower costs to purchase and to operate per-seat. The server, or a cluster of servers has the full weight of all the applications, services and data. By keeping a few servers busy and many thin clients lightly loaded we can expect easier system management, and lower costs as well as all the advantages of networked computing: central storage/backup and easier security.

Because the thin client is relatively passive and low-maintenance, but numerous, the entire system is simpler and easier to install and to operate. As the cost of hardware plunges and the cost of employing a technician, buying energy and disposing of waste rises, the advantages of thin clients grow. From the user's perspective the interaction is with monitor, keyboard and pointer changes little from using a thick client.

A single PC can usually power 3-6 or more thin clients. A more powerful PC or server can power up to 30 thin clients at a time. These are usually connected through a router and RJ-45 cables.

Thin clients are a great investment for schools and businesses who want to maximize the number of workstations they can purchase on a budget. A company called "ncomputing" has sold 1000's of access terminals to schools across the US, as well as other countries. A simple $70 unit could replace a computer in a school or business. It would also save a lot of power in the long run, due to the low power consumption of these units.

Thin Client Network. Users may log in on one server and run applications on multiple servers.

[edit] History

An IBM EXX Thin Client

What are now called thin clients were originally called "graphical terminals" when they first appeared, because they were a natural development of the text terminals that had gone before them. (Text terminals are essentially the ultimate thin client but are generally not classified as such - coming from an earlier computing era.)

X terminals were a relatively popular form of graphical terminal in the 1990s.

Late in the Windows NT 3.51 lifecycle, Citrix Systems approached Microsoft with an idea for a multi-user version of Windows similar to what had been done with Unix. Microsoft agreed to license the Windows NT 3.51 source code which Citrix then turned into a product called WinFrame; a version of NT 3.51 that allowed multiple users to run on the same server. Microsoft later licensed the technology back from Citrix and incorporated it into a special version of NT 4.0 (known as NT 4.0 TSE, or Terminal Server Edition) and then into all subsequent version of their Server operating systems. The code name for this Microsoft project was Hydra.

Terminal Services allows the operation of standard Windows software in a mainframe model - centralized computing vs. distributed computing. Users log onto the server using thin client hardware and the server creates a session in memory dedicated to that user. Any GUI commands that would normally be sent to a local graphics card are instead compressed and sent to the client. Likewise, user keyboard and mouse movements are sent back to the user's task running on the server.

It is likely that the term "thin client" started to be used instead of "graphical terminal" for the following reasons:

- When thin clients started to come back into vogue, fat clients had long been the norm in most environments. Many IT workers and managers used to working with fat clients such as PCs and Macs would have been unfamiliar with the term "graphical terminal".

- The term "thin client" is more descriptive and relevant than "graphical terminal", in an age in which *all* desktop computing devices have graphical capabilities.
- As a marketing term, it sounds short and snappy[*citations needed*] – and also, importantly, it made the technology sound innovative and technologically advanced, even though it was neither – X terminals had been acting as thin clients years before the term was widely used in the IT industry.
- "Thin Client" also reflects the fact that most of these devices leave out much of the hardware found in typical PCs, such as hard drive, cooling fan and much of the RAM.

[edit] Definitions

A thin client (or a lean client) is a network computer without a user writable long term storage device, which, in client/server applications, is designed to be especially small so that the bulk of the data processing occurs on the server. The embedded OS in a thin client is stored in a "flash drive", in a Disk on Module (DOM), or is downloaded over the network at boot-up. The embedded OS in a thin client usually uses some kind of write filter so that the OS and its configuration can only be changed by administrators.

Thin client (computing): A server-centric computing model in which the application software, data, and CPU power resides on a network server rather than on the client computer.

[edit] Application program

A thin client as an application program communicates with an application server and relies for most significant elements of its business logic on a separate piece of software, an application server, typically running on a host computer located nearby in a LAN or at a distance on a WAN or MAN.

A thin client does most of its processing on a central server with as little hardware and software as possible at the user's location, and as much as necessary at some centralized managed site.

Other definitions of thin versus thick/fat client application program try to draw the line at whether the deployment of the application requires the installation of additional software at the user site or not. Unfortunately, this is also arguable, since e.g., a browser used for a client application might be part of one client platform, but not the other. So on one platform no additional software installation is required, while another client platform requires it. The only objective definition would seem to be whether the boot image that is normally used to start the user's computer needs to be modified in any way before the client can be used: if not, then, the client is probably thin. Another criterion is related to the management of the thin client device or program. If it can be centrally managed, it is probably thin.

However, a great deal of software is today typically included in a base boot image, specifically to support various user applications, so that it need not be reinstalled on every computer. Often, a departmental boot image is prepared to include applications specific to a department.

[edit] User-interface device

A thin client as a device is designed to provide just those functions which are useful for user-interface programs. Often such devices do not include hard disk drives, which may become corrupted by the installation of misbehaved or incompatible software, but instead, in the interests of low maintenance cost and increased mean-time between failures (MTBF) the thin client device will use read-only storage such as a CD-ROM, Network Virtual Drive or flash memory.

Ideally the user will have only a screen, keyboard, a pointing device (if needed) and enough processing power to handle display and communications. Numerous companies develop and market these devices.

[edit] Device for running a thin client application program

A Gigabyte TA7 Thin client

"Thin client" has also been used as a marketing term for computer appliances designed to run thin client software. The NEC US110, IGEL Technology Universal Desktops, Wyse Winterms, Neoware's Appliances, Hewlett-Packard HP Compaq t-series, Chip PC Jack PC and Xtreme PC Series, SaaS style Nexterm,NEXterminal ,Sabertooth TC , ACP's ThinManager Ready Thin Clients, X terminal, ClearCube, Koolu, ThinCan or web kiosk might be considered thin clients in this sense.

A more recent concept in this genre is 'Ultra Thin Client' technology – which takes the 'thin' concept one step further by running the connection client software (Citrix, Windows Terminal Services, telnet etc) directly from the appliance's hardware. This is a marked difference to legacy thin-client hardware architecture which ran an operating system, often Windows CE or Linux between the hardware and connection client software.[2][3]

The latest concept is a 'Zero Client' which is a hardware only appliance that runs no software at all and directly connects the user to a virtualized desktop. This approach completely eliminates the need to manage thin client software of any kind. Pano Logic introduced such a device in 2007.

[edit] Software thin client

Many thin clients are software-only however, and run on standard PC hardware.

Examples of this type of software-only thin client include PXES Universal Linux Thin Client, Pilotlinux or Lan Core.[4] Knoppix is also actively pursuing this market, as is Thinstation (See also Puppy Linux). The SoThin Thin Client offers a Windows based solution and offers multiple connections to different platforms. Hopnetix is a diskless, network-booting thin client operating system that runs from a Windows server platform, similar to (and potentially licensed from) Neoware Image Manager software. Neoware Image Manager allows a diskless client to stream a complete system disk "on-demand" from a file server. In that perspective, it replaces the IDE or SCSI cable that connects a disk drive to a motherboard by a network protocol and a server that shares virtual disk drives to the client.

Another example on the Windows platform is the BeTwin program by the Thinsoft[5] company based in Hong Kong. The program, downloadable over the net, uses extra VGA/DVI graphics ports, or separate cards, in the host PC plus standard USB/PS2 connected keyboards, mice and sound systems to enable additional workstations and supports up to 5 user stations per PC. Each user has simultaneous and independent utilization of the operating system, installed programs and peripheral devices connected to the shared PC system. Running Windows XP on an adequate PC system plus use of dual head video cards

and the correct drivers, the software even supports two users per installed video card, additionally hardware efficient.

[edit] Examples of thin client and thick client usage

The advocates of both architectures tend to have contentious relationships. In practice, there seems to be little to choose between the two approaches for many applications. A few situations may clearly call for one or the other. Distributed computing projects such as the SETI@home project (whose whole point is to pass off computationally intensive analysis to a large collection of remote computers) are applications that require thick/fat clients. A classroom environment, with desk real-estate at a premium, large numbers of students in a single room, and the need to have every workstation run the exact same software in the exact same way, would definitely lend itself to a thin client solution. Several companies now sell thin-client laptops that can access internal resources through a Virtual private network so the connection between client and server passes through an encrypted tunnel. This can allow mobile workers to access security-sensitive databases with less risk of lost or compromised data should the laptop be lost or stolen since it has no local storage.

[edit] Industrial thin client applications

Since 2006 there has been a growing interest in using Thin Client technology in hazardous areas, such as oil & gas exploration, military mobile use to monitor gen sets and mobile missile installations, and in industry in Zone 1 areas where hardened industrial computers can be prohibitively expensive. Thin Client hardware is easier to seal against environmental hazards and contamination, and can sometimes withstand a wider temperature and vibration level, due to simplified components and lack of moving parts, such as hard drives and cooling fans.

Thin Clients are also a natural choice for operator interface displays in manufacturing areas where a rugged tamper proof solution is required, and where limited "stateless" applications such as Human Machine Interface (HMI) and Supervisory Control and Data Acquisition (SCADA) software is the standard. "PC" on a plant floor usually stands for "Problem Center".[*citation needed*]

Flash memory is a critical component of these Thin Client terminals. With sufficient Flash (solid-state) memory and the ability to install user firmware and PCI cards, a Thin Client can perform most of the functions that a user would want to perform at these locations, typically called "point-of-attack" or "rig floor". The Thin Client can be equipped with a sealed membrane keypad for operator interface, and I/O (input/output) connectors that allow the Thin Client to be hooked up to a remote location and perform data recording of processes or history of connected device readings, that could be anything from radar to pumps to drilling apparatus.

Industrial Thin Clients typically operate on embedded software systems such as Linux, Windows CE.net, or Windows XP Embedded Ethernet. Industrial clients typically prefer an OS (operating system) environment in which the Thin Client can easily interface with proprietary firmware (software embedded in memory on a PCB card), and communications often interface with a hosting central computer in a safe environment control room. Ethernet protocol is a preferred network means of interfacing the computers, Thin Clients, and devices being monitored and controlled.

ProPanel Industrial Thin Client

Some Thin Client solutions (such as ACP's ThinManager Ready Thin Clients) are tightly coupled with specialized management software that enhances the basic features offered by server operating systems. These enhancements include, but are not limited to, failover to backup servers, client device independence, redundant Ethernet, multiple sessions on a single client and automatic client configuration.

Fiber networks and wireless interface are also popular to enhance versatility in location and installation of Thin Clients. Encryption of data is usually a priority, and wireless options need to rely on a fast Ethernet radio transmitter using 11Mbit/s 2.4 GHz 802.11 output, otherwise data transmission speed and security can be compromised.

Ruggedized Thin Client products therefore enable easy-to-employ industry standard network creation and control at hazardous area zones for less cost and with less risk of failure than full computer systems. In fact, in the first quarter of 2007, mandates have been created by the US Armed Forces to look at Thin Client solutions in all field applications. The military is primarily interested in Thin Client technology in the field due to its improved cost control, more robust construction, less vulnerability to failure and security breaches, lesser weight and greater mobility, and lower incidence of OS failures.

[edit] Advantages of thin clients

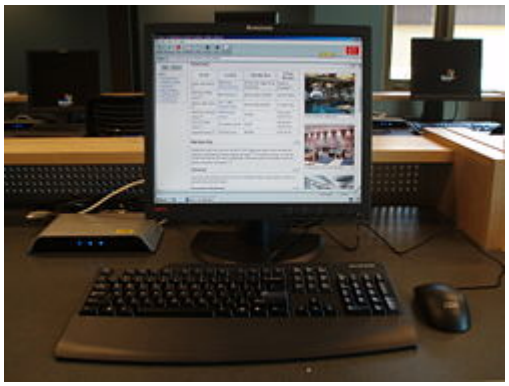

A public thin-client computer terminal inside a public library.

Obviously, boot image control is much simpler when only thin clients are used – typically a single boot image can accommodate a very wide range of user needs, and be managed centrally, resulting in:

- **Lower IT administration costs**. Thin clients are managed almost entirely at the server. The hardware has fewer points of failure and the client is simpler (and often lacks permanent storage), providing protection from malware.
- **Easier to secure**. Thin clients can be designed so that no application data ever resides on the client (just whatever is displayed), centralizing malware protection and reducing the risks of physical data theft.

- **Enhanced data security**. Should a thin-client device suffer serious mishap or industrial accident, no data will be lost, as it resides on the terminal server and not the point-of-operation device.
- **Lower hardware costs**. Thin client hardware is generally cheaper because it does not contain a disk, application memory, or a powerful processor. They also generally have a longer period before requiring an upgrade or becoming obsolete. There are fewer moving parts and one upgrades the server and network instead because the limitation on performance is the display resolution which has a very long life cycle. Many thick clients are replaced after 3 years to avoid failures of hardware in service and to use the latest software while thin clients can do the same, well-defined task of displaying images for 10 years. The total hardware requirements for a thin client system (including both servers and clients) are usually much lower compared to a system with fat clients. One reason for this is that the hardware is better utilized. A CPU in a fat workstation is idle most of the time. With thin clients, CPU cycles are shared. If several users are running the same application, it only needs to be loaded into RAM once with a central server (if the application is written to support this capability). With fat clients, each workstation must have its own copy of the program in memory.
- **Less energy consumption**. Dedicated thin client hardware has much lower energy consumption than typical thick client PCs. This not only reduces energy costs but may mean that in some cases air-conditioning systems are not required or need not be upgraded which can be a significant cost saving and contribute to achieving energy saving targets. However, more powerful servers and communications are required.
- **Easier hardware failure management**. If a thin client fails, a replacement can simply be swapped in while the client is repaired; the user is not inconvenienced because their data is not on the client.
- **Worth less to most thieves**. Thin client hardware, whether dedicated or simply older hardware that has been repurposed via cascading, is less useful outside a client-server environment. Burglars interested in computer equipment may have a much harder time fencing thin client hardware.
- **Operable in Hostile Environments**. Most thin clients have no moving parts so can be used in dusty environments without the worry of PC fans clogging up and overheating and burning out the PC.
- **Less network bandwidth**. Since terminal servers typically reside on the same high-speed network backbone as file servers, most network traffic is confined to the server room. In a fat client environment if you open a 10MB document that's 10MB transferred from the file server to your PC. When you save it that's another 10MB from your PC to the server. When you print it the same happens again — another 10MB over the network to your print server and then 10MB onward to the printer. This is highly inefficient. In a thin client environment only mouse movements, keystrokes and screen updates are transmitted from/to the end user. Over efficient protocols such as ICA or NX this can consume as little as 5 kbit/s bandwidth.[*citation needed*][*dubious* – *discuss*] This statement makes some very heavy assumptions about the operating environment, though.
- **More efficient use of computing resources**. A typical thick-client will be specified to cope with the maximum load the user needs, which can be inefficient at times when it is not used. In contrast, thin clients only use the exact amount of computing resources required by the current task – in a large network, there is a high probability the load from each user will fluctuate in a different cycle to that of another user (i.e. the peaks of one will more than likely correspond, time-wise, to the troughs of another. This is a natural result of the additive effect of many random, independent loads. The total load will be normally distributed about a mean and not the sum of the maximum possible loads. see Central Limit Theorem The resulting mean load is *n* times the mean independent load with a standard deviation $\frac{1}{\sqrt{n}}$ times the standard deviation of the independent loads. The probability of a total load more than a few standard deviations above the mean load is extremely small.

- **Simple hardware upgrade path**. If the peak resource usage is above a pre-defined limit, it is a relatively simple process to add another component to a server rack (be it power, processing, storage), boosting resources to exactly the amount required. The existing units can continue to serve alongside the new, whereas a thick client model requires an entire desktop unit be replaced, resulting in down-time for the user, and the problem of disposing of the old unit.
- **Lower noise**. The aforementioned removal of fans reduces the noise produced by the unit. This can create a more pleasant and productive working environment.
- **Less wasted hardware**. Computer hardware contains heavy metals and plastics and requires energy and resources to create. Thin clients can remain in service longer and ultimately produce less surplus computer hardware than an equivalent thick client installation because they can be made with no moving parts. Computer Fans and disk storage (used for cooling and storage in thick clients) have mean times before failures of many thousands of hours but the transistors and conductors in the thin client have mean times before failure of millions of hours [1]. A thick client is considered old after one or two cycles of Moore's Law to keep up with increasing software bloat but a thin client is asked to do the same simple job year after year. A thin client, on the other hand will be replaced only when it lacks some feature deemed essential. With audio, video, and USB, thin clients have changed little in 15 years, being essentially, stripped-down PCs.

[edit] Advantages of thick clients

- **Fewer server requirements**. A thick client server does not require as high a level of performance as a thin client server (since the thick clients themselves do much of the application processing). This may result in cheaper servers although in practice many thin client servers are actually equivalent to file servers in specifications but with additional memory.
- **Better multimedia performance**. Thick clients have advantages in multimedia-rich applications that would be bandwidth intensive if fully served. For example, thick clients are well suited for video editing and video gaming.
- **More flexibility**. On some operating systems (such as Microsoft Windows) software products are designed for personal computers that have their own local resources. Trying to run this software in a thin client environment can be difficult or impossible, especially for applications that have many shared objects or libraries that are accessed frequently.
- **Better peripheral support**. Thin clients are typically very small, sealed boxes with no possibility for internal expansion, and limited or non-existent possibility for external expansion. Even if for example, a USB device can be physically attached to a thin client, the thin client's software might not support peripherals beyond the basic input and output devices - for example, it may not be compatible with graphics tablets, digital cameras or scanners.
- **Suitable for poor network connections**. Thin clients can be unusually slow, or very frustrating to use, over a high latency network connection. Moreover, they do not work at all when the network is down. It may be possible to work offline with a thick client, although the network oriented manner in which many people work today means that thick client usage can still be curtailed if the network is down.
- **Easier to repurpose**. May be used in thin client applications when the hardware becomes obsolete for thick client use. Because it is standard and can operate autonomously, thick client hardware is easier to resell or donate when it must be retired.

[edit] Thin client management software

Most moderate to large Thin Client deployments use some type of management software. Companies such as Citrix (XenApp), IGEL Technology (Remote Management Suite), Wyse (Rapport), 2X (ThinClientServer), ACP (ThinManager), ThinSoftInc and Symantec (Altiris Deployment Solution) have products that vary in cost and functionality. These products all enhance the basic server operating system,

adding features such as Server Load Balancing, Application Publishing, Shadowing, Local Storage Options and Session Reconnection.

[edit] Client/server protocols

ALP
: Appliance Link Protocol is a secure bitmap-based network protocol used by Sun with its Sun Rays.

XML over HTTP
: Protocol used by the AJAX model of web applications.

X11
: Networking and display protocol available on essentially all Unix variants.

X11 over Secure Shell
: Tunneling an X11 session through an ssh session to provide authentication and encryption.

NX technology
: Compresses and caches the X11 protocol for better performance.

VNC
: Allows for (virtual) desktop sharing.

Citrix ICA
: with Citrix Presentation Server

RDP
: The default remote desktop access mechanism for MS-Windows

RGS
: A client-server software solution developed by HP, to enable remote access to high-performance workstations from a thin-client machine.

HTML over HTTP
: Used by a myriad of web applications.