




Linux Thin Client Networks Design and Deployment

A quick guide for System Administrators

David Richards



Chapter No. 5 "Considering the Network"

In this package, you will find:

A Biography of the author of the book

A preview chapter from the book, Chapter NO. 5 "Considering the Network"

A synopsis of the book's content

Information on where to buy this book

About the Author

David Richards is a System Administrator for the City of Largo, Florida. He has been exposed to computer technologies since the day he got his first home computer in the early 1980s. After graduating from college in 1986, he was employed in the manufacturing, distributing, and printing industries. 1992 was the first year that he entered the City's employment and has been working with Unix, Linux, and thin clients. He promotes thin clients and open source technology, and enjoys the challenges in their deployment. He is often found in the GNOME IRC channels debugging software and interacting with the developers.

For More Information:

www.PacktPub.com/linux-thin-client-networks-open-source/book

Linux Thin Client Networks Design and Deployment:

A quick guide for System Administrators

It is with interest that the author has an eye on technology when visiting other organizations and also as part of normal daily life. A visit to the local home improvement or video store will show how expensive client/server technology has been deployed, where thin clients would easily meet their needs. It's also interesting to the author that a visit to a major computer store chain demonstrates that solid-state computing is known for its stability and reliability. This entire store is devoted to personal computers and networking, and all of their point of sales and inventory systems are on dumb terminals!

Thought has been given to this book in terms of creating something that is well rounded, and meets the needs of small and large organizations. Computer technology is changing all the time, and attempts were made to keep the contents of this book relevant for as long as possible. Items of security are best suited for the System Administrator or integrator of a solution, and will not be discussed greatly in the chapters. The information covered will give you enough knowledge to understand how the technology works, make decisions about deployment, and then implement a stable work environment.

What This Book Covers

Chapter 1: Overview of Thin Clients will give you an overview of what exactly is a thin client, and the different types of models that are available.

Chapter 2: The Types of Thin Clients will give you a sample of multiple types of thin clients. Once you have configured your server, you will have to make decisions about which types of hardware to deploy.

Chapter 3: An Analysis of Costs will identify key areas to review when considering the financial impact of your thin client plan. This includes hardware acquisition and also staffing costs.

Chapter 4: The People Issues will address what might be the hardest part of your deployment: People. Some people are passionate about their software and others are challenged with any work flow changes. It's important to address them as much as possible before, during and after deployment.

Chapter 5: Considering the Network will review the network required to run thin clients. Because of the simplicity of the computing deployment, your network too is simplified.

For More Information:

www.PacktPub.com/linux-thin-client-networks-open-source/book

Chapter 6: Implementing the Server will cover the steps necessary to design a server for the number of users in your deployment. Also covered will be steps to allow thin clients to log into and run a desktop environment.

Chapter 7: Implementing the User Software will provide ideas for software packages that run on Linux, along with their suitability to run over the network to thin clients.

Chapter 8: Implementing the Thin Clients reviews the process of considering the operating system to deploy on the devices. Also covered is the interaction with USB devices and speakers.

Chapter 9: Support covers three aspects of support. The first is supporting your users, the second is support within your IT staff and the third is support from software vendors.

Appendix A: This lists out the URLs of the various projects and hardware mentioned in this book.

Appendix B: This takes you through the installation of OpenSuse.

For More Information:

www.PacktPub.com/linux-thin-client-networks-open-source/book

5

Considering the Network

Much information already exists concerning methods for deploying networks and hardware. We shall try to clarify differences between using a network with personal computers and thin clients with the help of anecdotes. Certain designs have also proven to be very stable and provide the best possible solution.

Primary Network

Your first thought might be that your current network will work fine with thin clients and that is entirely possible. But your network might be something that has grown through the years and is not that well designed. Your implementation of thin clients then might be a good time to review the design and make upgrades as needed.

Personal Computers versus Thin Clients

Based on conversations with some hardware vendors, it's clear that most of the testing is done with the expectation that personal computers will be deployed.

For More Information:

www.PacktPub.com/linux-thin-client-networks-open-source/book

The biggest difference is in how the two platforms use the network. When running a personal computer, often software applications are stored on network servers. When you activate an icon, the network pushes the executable down to your PC. Once downloaded into memory, the application runs and then very little interaction takes place until you save a file. Or in other cases, the executables are on the local PC, and network activity is not used until files are saved. If an executable takes a few seconds longer to download, you won't really notice it when using a personal computer. Some networking devices seem better designed for efficiency of download instead of being designed for the smaller and more plentiful packets of network computing. When you activate a software application on a thin client, the presentation of the user interface is pushed to you from the server, and then all keystrokes and mouse activity are transmitted back and forth to the server in real time. The network needs to be very fast, have low latency, and be configured to pass packets immediately to the servers.

Network Design

For implementing your network, the network backbone should be Gigabit if possible. Obviously if your solution is for only a small number of users, this might not be required. Ideally fibre optic lines are then run to each of the wiring closets, and each switch should have it's own line. It is advisable to avoid daisy-chaining the switches together in order to avoid any kind of contention between them. The servers are all plugged into the backbone at Gigabit as well. If a server is required away from a centralized computer room, then it is better to run a separate line instead of plugging it into a switch that will be shared with thin clients. It's important to keep the data paths solidly designed so that all of your real-time interaction will not be delayed.

X windows, RDP, or Citrix are used to display the user presentation. This means that the software is running on the server, but the image of that software is transmitted over the network. It's important that a strong network exists or repaints of windows will be slower and feel sluggish. This issue will cause people problems, with perceptions that a personal computer can run software faster than a network. A correctly designed network will provide excellent response time and the user community should not even see a difference.

Font servers are used to distribute fonts to users. A font server is just a process or application that runs on the server. When a user requests a font, it's sent over the network to the thin client and made available to them immediately. The strength of this design is that all your employees will have the same fonts and while sharing documents, they will render exactly the same way no matter from where you log into the network. Anyone that has shipped documents between personal computers with different fonts, will greatly appreciate this design. When the network is configured correctly, font download and interaction is immediate and undetected by the user community.

NFS mounts are used to connect disk drives between Linux servers. This allows applications to share data between the various servers on your network. Response time needs to be excellent to provide very fast file saves and retrievals, at the same time avoiding applications that lock or timeout while trying to interact with files.

A review of the possible network problems is provided in the following table:

Networking Issue	Symptom(s)
X windows	Slow repaints of user interfaces
RDP	Lockups
Citrix	Disconnections
Font Servers	Slow repaints
Thin Client	Wrong fonts in applications
	Slow keyboard response
	Disconnections
NFS Mounts	Slow repaints
	Lockups
	Slow response in saving files

Remote Sites

It would be wonderful if Gigabit could be run to all of your facilities. But the truth is that often you are not able to deploy that speed, because of cost or physical locations of buildings. Once networking speeds get below 100 Megabit, you will no longer be able to use native X windows and must consider deploying products that compress presentation data. Microsoft RDP will do this, along with Citrix Metaframe, tight-VNC, and NX/Nomachine.

Most of the products the author has tested that compress data seem to become usable around 100K of speed. Dialup connections will work, but repaints will be tedious and not very efficient. A good formula to use is to multiply the number of concurrent users at each remote site by 100 to get a rough estimate of bandwidth required. Using this formula, 3 concurrent users would require roughly 300K of bandwidth. Remember too, that

very possibly print jobs will be running on the same circuit, which will consume bandwidth as well. The user community will perceive 'slowness' mostly in the user presentation itself. Print jobs that take a bit longer are not normally noticed. So, one might consider running two circuits to remote sites and splitting the user sessions on one, and the print jobs on the other. That way if massive print jobs are sent, the users won't notice and can continue working. It should be noted as well that some bandwidth management products such as Citrix have designed their software to support printer connectivity that is also compressed.

The most important issue with remote sites is stability and uptime. When you centralize all of your software, it's critical that the network be available or users will not be able to log into the servers and do their job. Many people do not care how it works, just that it's reliable. Consider all of your options such as T1 connections, DSL, and cable modems and then select the solution that seems like the best fit. One effective method is to create a list of all available networking methods, and then create a chart that clearly spells out the features and speed available within each category. As the line becomes cheaper, it normally becomes slower. The decision makers need to understand that at a certain point presentation and software application speed will start to degrade. It is also important to obtain from the vendor exact service levels for each of the connection methods. Commercial and business lines often will guarantee a minimum amount of bandwidth. Regular home-use circuits often are rated for 'burst rates' and run considerably slower than the specified rate.

Some remote users will be using wireless connections. They too will require an application that will perform compression, and should be considered as well. Cellular wireless broadband is providing plenty of bandwidth these days to run with centralized computers.

Thin Client Network Connections

Once you have performed the steps as outlined previously, you will need to finalize your design for the thin clients themselves. If your current deployment is only capable of 10 Megabit, you will definitely want to upgrade the wiring and move to a minimum of 100 Megabit. 100 Megabit provides plenty of bandwidth for very crisp response time from the servers, and mouse and user interface response is excellent. At the time of this writing, running Gigabit to the desktop has not been tested by the author; but if your thin clients support that speed that is an excellent buffer and should provide even greater capabilities. If users are using devices such as laptops, always encourage them to use a wired connection in the office. In the case of X windows, it will avoid having to use a licensed bandwidth compression interface.

At the time of this writing, Gigabit connections to the desktop are becoming more and more commonplace. If part of your deployment is a redesign of the network, run the highest-rated wiring that you are allowed. The author is anticipating that Gigabit thin clients will become available very shortly. More and more software is making use of the 3-dimensional capabilities of thin client video cards, and each step in this direction requires additional bandwidth.

Testing the Network

Anyone that has supported users knows that often they will discover things that the technical staff never anticipated. It's important to turn that into a positive, and isolate as many problems as possible before final deployment. It's effective to place a few thin clients at each of your sites, and on each of the networking technologies that you have selected and perform their regular day-to-day duties. Some types of connections such as X windows sessions are not stateless, and will drop if the network under-performs. If you are considering a new vendor for networking hardware, they should allow you to install demo devices and test them

on your infrastructure. Be mindful that sales people sometimes over sell their products or don't understand exactly your design goals, so a real-world test with their hardware before major purchases is always a good idea.

Summary

In this chapter we see that though the complexity of networking cannot be stated strongly enough, it is important to design a rock-solid and stable network before your deployment begins. Follow standard methods and designs and work with your hardware vendors to make the best possible use of their equipment. Once you plug the thin client in the wall, you will be excited at the things to come and will be ready to configure the application servers.

Where to buy this book

You can buy Linux Thin Client Networks Design and Deployment from the Packt Publishing website: <http://www.packtpub.com/linux-thin-client-networks-open-source/book>.

Free shipping to the US, UK, Europe, Australia, New Zealand and India.

Alternatively, you can buy the book from Amazon, BN.com, Computer Manuals and most internet book retailers.



www.PacktPub.com

For More Information:

www.PacktPub.com/linux-thin-client-networks-open-source/book