May 2005

Information Technology and Systems - II: Server Administration Networks

Adam Victor Reed

California State University, Los Angeles, areed2@calstatela.edu

Follow this and additional works at: <https://aisel.aisnet.org/cais>

Recommended Citation

Reed, Adam Victor (2005) "Information Technology and Systems - II: Server Administration Networks," *Communications of the Association for Information Systems*: Vol. 15 , Article 35.

DOI: 10.17705/1CAIS.01535

Available at: <https://aisel.aisnet.org/cais/vol15/iss1/35>

This material is brought to you by the AIS Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in Communications of the Association for Information Systems by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

INFORMATION TECHNOLOGY AND SYSTEMS – II: SERVER ADMINISTRATION NETWORKS

Adam Reed
Computer Information Systems
California State University, Los Angeles
areed2@calstatela.edu

ABSTRACT

A majority of IS graduates (56% in one recent survey] are involved in server administration, network administration and IS security work. An important recent innovation in these areas is the deployment of separate networks dedicated to server administration and related tasks, combining the cost and productivity advantages of remote administration with risk levels comparable to console-based administrative access. Remote server administration is a previously undocumented artisanal tradition that evolved in scientific and technical network environments, and is now becoming applicable to an increasing range of business networks. This tutorial article provides an overview of current server administration network architectures, and of the software, workstation, and user interface technologies associated with remote server administration.

Keywords: remote administration, productivity, security, networks, workstations.

I. INTRODUCTION

Network-based servers are central to current operational praxis in business and organizational information systems. The Department of Computer Information Systems at California State University, Los Angeles recently surveyed graduates who received degrees in Information Systems from our Department in the last six years. Of the 131 recent graduates who responded to our on-line survey, 57 (44%) functioned as Windows server administrators, and 31 (24%) had direct report subordinates whose tasks included Windows server administration. The corresponding numbers were 11 (8%) and 8 (6%) for Unix/Linux server administration; 55 (42%) and 30 (23%) for network administration; and 35 (27%) and 16 (12%) for IS security administration. In total, 73 (56%) of our responding alumni performed or supervised server, network, or security administration. To serve the needs of students who will work in these areas, IS curricula should incorporate instruction relevant to their future assignments. The present tutorial is intended to introduce an important current development in server, network, and security administration: the replacement of console-based server administration with remote administration using dedicated server administration networks. Remote server administration has been, prior to the present article, a largely undocumented artisanal tradition. This article documents server administration network concepts and technologies, both for working IS professionals and for IS faculty who will be preparing their students for the current state of professional practice; and includes a glossary of relevant terms.

II. A HISTORY OF NETWORK PARTITIONING

The deployment of dedicated server administration networks is the culmination of a longer-range technological trend toward separate networks for different functions. In the two decades before 2005, organizational local area network architectures moved from the single network model - still depicted today in many data communication textbooks - to multiple networks partitioned from one another by intervening servers and firewalls. As recently as 1985, most internal organizational networks were segments of the worldwide Internet; internal computers bore Internet Protocol addresses drawn from the address space of that worldwide Internet. Around 1985, a university or corporate network looked as shown in Figure 1.

Figure 1. A Typical Organizational Data Network Around 1985.

Improvements in the understanding of the risks of an unfiltered connection to the Internet subsequently resulted in the addition of packet filter firewalls that implemented address-based and port-based filtering rules in the organization's external gateway router. To provide a higher level of protection to internal clients and servers that did not need to face the external Internet, a stronger, typically stateful firewall¹ was interposed between internal servers and clients and the outer network segment, called a "demilitarized zone" or DMZ², that included the network interfaces of Internet-facing servers. With the deployment of separate networks on the internal side of the stateful firewall, it became necessary to use servers with interfaces to both networks simultaneously. Thus, the box labeled "Firewall and Boundary Servers" in Figure 2, below, represents, in addition to a stateful firewall, separate inbound and outbound SMTP [Postel, 1982] servers, any number of protocol-specific proxy servers for HTTP [Fielding et al 1999], HTTPS [Rescorla, 2000], SSH [Sommerfeld, 2003] and other protocols, VPN [Wilder et al 2004] tunnel servers, and any other servers that need to be interposed between the world-wide internet and the organization's internal client network. As a result of demand for machines capable of deployment at the interface between internal and external networks, server hardware manufacturers now provide two or more separate network interfaces as a standard feature of many server machines. By 1995, a typical organization's networks were organized as follows:

¹ A *stateful* firewall is a firewall that monitors the establishment of authorized TCP connections, and stops inbound packets containing TCP segments that do not belong to previously established authorized TCP connections.

² The use of the "demilitarized zone" or "DMZ" designation for the network segment between the external firewall and outward-facing servers, has been standard in the industry since the early 1990s; it was popularized by Cheswick and Bellovin [1994].

Figure 2. A Typical Organizational Data Network Around 1995.

In the following years, the Internet experienced increasing congestion of the worldwide IP (version 4) address space, and a resulting scarcity of IP addresses. The solution to this problem was network address translation (NAT) [Egevang and Francis 1994.] With NAT, internal machines that do not need to receive incoming external connections use IP addresses drawn from a separate, private IP address space (usually a subset of 192.168.*.*) The vectors of IP addresses and TCP port numbers of outbound connections from those (possibly hundreds of) internal machines are translated to TCP port numbers on a single externally visible IP address. The network address translator enabled a separation between internal networks, using the internal IP address space, and the external, worldwide Internet – with the network address translator interposed between them, replacing previously deployed stateful firewalls.

The availability of machines with multiple network interfaces encouraged the deployment of separate, dedicated networks in any situation in which this would benefit efficiency or security. For example, back-end servers, such as corporate database servers, are now typically connected with separate networks to the organization's internal and external servers. Those separate networks, shown in Figure 3, eliminate the security problems that resulted from allowing connectivity between the relatively open internal client network, and the highly sensitive back-end servers. Permitted types of access to restricted resources on the sensitive and highly protected back-end servers will differ, depending on the network interface originating the request. Typically, internal servers have greater privileges than the more readily compromised DMZ-facing external servers. Thus around 2001 the typical corporate data network looked like this:

Figure 3. Typical Corporate Data Network Architecture Around 2001.

II. RECENT DEVELOPMENTS

SERVERS DEDICATED TO A SINGLE SERVICE

In recent years, radically lower prices of server hardware, the availability of firewall appliances, and thus the possibility of realizing security and efficiency gains from dedicating each server computer to one specific server application only, have driven a universal transition to highly distributed server network architectures. The tradition of running corporate server software in small and medium-size business organizations on one or a few computers, each of which ran several (or many) services simultaneously, has been replaced by the current practice of dedicating one or more separate computer(s) to every network-accessible service.

One motivation for this change is to reduce vulnerability to attacks. Most attacks against servers exploit buffer overflows and other defects in the software that implements a specific service. Once an attack succeeds, the attacker can gain administrative control of the server operating system, and thus compromise the data and the integrity of all the services running on the subverted machine. Even the most secure database, for example, may be compromised through a software defect in a publicly accessible mail or web server running on the same machine. Dedicating a separate machine to each service substantially reduces the vulnerability of every affected service.

The performance of service software is affected not just by the speed of a machine's processors, but also by the need, if any, to transfer the content of the memory belonging to a given software process from RAM to disk, and back, when the RAM is temporarily allocated to a different process. Dedicating a separate machine to each service substantially reduces, and often eliminates, inter-process contention for the machine's physical RAM. When data swapping between disk and RAM is minimized, the service process to which a given machine is dedicated can run with optimal efficiency.

SEPARATE SERVER ADMINISTRATION NETWORKS

The strategic change from the use of one or a few server machines to dozens, and often hundreds of servers, is driving the replacement of console-based administration with network-based administration using dedicated *server administration networks*. In console-based administration, the configuration and maintenance utilities of server applications were run on physical console equipment - a keyboard, a video display, and a mouse (KVM) - connected to each machine. Moderate increases in the number of server machines were handled with KVM switches, which switch a single physical KVM console among the physical KVM interfaces of several server computers. However, KVM-switching technology becomes unwieldy (and expensive) when applied to networks of more than a few dozen servers. The alternative is to use a separate administrative network.

Figure 4 shows the topology of the contemporary (2005) corporate network architecture. Each box in this figure represents anything from a single computer to dozens, hundreds, or even thousands of computer systems. For example, as noted above, the box labeled "Firewall and Boundary Servers" represents one or more network firewalls, separate inbound and outbound electronic mail (SMTP) servers, any number of protocol-specific proxy servers (for HTTP, HTTPS, SSH and other protocols,) VPN tunnel servers, and any other servers that need to be interposed between the world-wide internet and the organization's internal client network.

The transition to remote server and router administration through a separate administrative network is an application of the principle that

the risk of unauthorized access to information resources and processes is minimized by running separate networks at different levels of trust.

Figure 4. Typical 2005 Network Architecture with server Administration Network.

In modern operating systems, every software process can trace itself to the interface (or system daemon) by which it was spawned and by which it is controlled. This arrangement permits each administrative application on a server platform to be configured so that it can only be controlled from one specific network interface. Similarly, access to resources such as database records and fields may be restricted on the basis of the network interface from which the request for the resource originates. Permitted types of access will differ, depending on the network interface originating the request.

Except for some small computers built exclusively for client use or for specific control applications, all current hardware platforms that are used to deploy server applications support multiple network interfaces. Similarly, the workstations used by server and network administrators are often based on server hardware and also support multiple network interfaces. Each category of machines shown in Figure 4, other than internal clients, is configured with network interfaces to multiple networks. Administration workstations are connected both to the internal client network for work components other than server administration, and to the administration network. Note that the category of administration workstations connected to the administration network includes workstations for all functions that require a special, trusted level of access to corporate servers. For example, a webmaster authorized to modify the content of the company website will do so using the administration network.

TRAFFIC ON THE SERVER ADMINISTRATION NETWORK

The traffic on the server administration network consists of five broad categories.

1. The administrative network is used for secure transfer of sensitive files, such as web content and server configuration files, from administrative workstations to servers.
2. The network carries traffic of the Simple Network Management protocol (SNMP, [Case et al. 1990]), which is used to manage routers and other network devices, between the

SNMP software agent on each managed component, and the central SNMP client on the network administrator's workstation.

3. Some simple network components and server applications are administered through HTTP or HTTPS World Wide Web protocols, with an administrative web server listening (sometimes on a non-standard TCP port) for requests from a standard browser running on the administrative workstation. For sufficiently simple tasks, the HTTP/HTTPS protocols provide a platform-independent server administration mechanism that does not require add-on software for either family of platforms.
4. The server administration network may carry remote backups of server file systems, so that backup devices need not be attached to any machine other than the workstation of the administrator responsible for backups. All backups can be performed to that workstation over the server administration network.
5. The most critical category of traffic over the server administration network distributes the user interfaces, of the administrative software running on the different servers, for operation through the display and input hardware of the administrative workstation.

This praxis of remote, network-based system administration entered the organizational and business computing environments only recently, as a result of the deployment of separate server administration networks. In its previous history, remote system administration was an artisanal tradition that evolved among administrators of research, academic, and technical computing systems. In the last few years this artisanal tradition spread, for reasons given above, to corporate IT networks of enterprises with research and technical systems, and employed system administrators in support of research and other technical activities. In 2005, server administration networks are increasingly deployed throughout the larger organizational IS community, in the hope that this will lead to more secure configuration and administration of organizational networks in general, beyond those supporting predominantly technical organizations. The most demanding aspect of the successful deployment of dedicated server administration networks is the implementation of remote desktop technologies and protocols.

IV. REMOTE DESKTOP TECHNOLOGIES

Server administration through the administration network depends on protocols that permit the user interface, of administrative applications running on a server machine, to be presented through the user interface hardware (keyboard, video display and mouse) of the administrative workstation. The traditional remote application and remote desktop protocols for Unix and Linux platforms are the X Window System [Scheifler and Gettys 1992] and XDMCP [Packard 2004, Braca 1991, Smith 2002 pp. 348-362]. On Microsoft platforms, functionality equivalent to X and XDMCP is provided through a proprietary Microsoft protocol, the Remote Desktop Protocol (RDP). The application end of one of these protocol sets is built into most server administration applications on the respective families of platforms. The user interface ends of the protocols are part of each workstation platform software distribution. Optionally, add-ons can be installed on each administrative workstation platform for remote administration of applications running on server platforms of the other family.

THE X WINDOW SYSTEM

The X Window System is the standard user interface system on Unix and Linux platforms. The architecture of Unix and Linux operating systems is based on the principle of partitioning functionality among self-contained modules (objects, subroutines, processes, commands, applications, kernels) each of which performs a single, specific function, or a very closely integrated system of functions. These modules are connected to each other with publicly documented interfaces. This principle implies that any module, including the kernel of the operating system, may be redesigned or replaced without affecting the code or the functionality of other modules. Since the origin of Unix in the early 1970s, every module has been replaced

several times, independently, with improved versions, while continuing to interoperate seamlessly and transparently with other modules.

Each of the publicly documented interfaces among the component modules of a Unix/Linux system implements a communication protocol. These protocols are generally independent of whether the communicating components are running on the same machine, and under the same kernel, or communicating with each other over a network. Thus, for example, a logging process need not run on the same machine as the kernel or application it serves, and may simultaneously serve kernels and applications on many machines over the network.

In line with the Unix/Linux architecture, the user interface of Unix and Linux systems is separate from both the kernel and the application processes. The given user's interface hardware - keyboard, pointer, and one or several display screens - is controlled by a component process called the X Window Server. Applications connect to the X Window Server as clients. The protocol of the X Window System is completely network-transparent, so that connected applications are equal, whether they are running on the same computer as the X Window Server to which they are connected, or connecting to it through a network from another system.

The user interface of the X Window Server itself is architecturally just another client process, called a "window manager". Since the window manager uses the same protocol to connect with the X Window Server as any other client, it too may run on any machine in the network. Thus the X Window System provides Unix and Linux machines with a built-in ability to control remote applications through a user interface displayed on the local desktop, or to import the entire user interface of another computer in one's network to one's local workstation. On Unix networks, the administrator of multiple servers was always able to use a single workstation to interact with the native administrative interface of any server. There was never a need to co-locate the servers, or to run keyboard, mouse and video cables from the servers to the administrator's desk. Even servers located remotely, on other continents if needed, could be fully administered, through their native user interfaces, from the administrator's desktop workstation. It is this capability that gave rise to the artisanal tradition of remote administration in organizations using Unix and Linux operating systems.

WORKSTATION DISPLAY MANAGEMENT AND XDMCP

On a typical Linux or Unix desktop computer with an X Window System display, when the display is not in use the operating system runs an application called a Display Manager. The usual function of the Display Manager is to let the user log into the local computer. The manager provides input fields for this purpose. Once the user has typed her account name and password into these fields, the Display Manager brings up the user's desktop on the local machine.

In addition to local login, Display Managers also support buttons which pop up menus that can perform some functions without the user logging in. On display managers designed for typical desktop client machines, these buttons and menus may permit the user to, among other things:

- Shut down or reboot the desktop machine
- Choose a language for the next login
- Choose the type of session that the Display Manager will start:
Use the command line or a safe, minimal desktop, CDE (Common Desktop Environment [The Open Group 1998], Gnome (GNU network Object Model Environment) [Warkus 2004], KDE (K Desktop Environment) [The KDE Team 2001]

These desktop environments provide different Window Managers, different startup and customized configuration capabilities, different styles, and capabilities and features that the user may adapt to his or her personal cognitive style and experience for maximum personal productivity.

The Display Manager of a workstation designed for network, server, and service administrators differs from a Display Manager designed for a simple desktop client in that it also provides an interface to XDMCP, the X Display Management Control Protocol, for remote access to the native user interface of any appropriately configured Unix or Linux-based server on the network. The Sun Microsystems Solaris Display Manager, for example, includes in its "Options" button menu a "Remote Login" sub-menu, with choices to "Enter Host Name" or "Choose Host From List." The "Remote Login" submenu invokes the capabilities of XDMCP.

If the user selects "Enter Host Name" from the XDMCP menu and specifies a name or an address on the network, the local display manager connects to the remote host – in this case, the server or router to be administered – through the administrative network. If the remote host is running the X Display Management (xdm) network service on its interface to the administrative network, the local Display Manager uses XDMCP to hand the control of its display to the Display Manager of the remote host. Since the Display Manager connects to the X Window System Server as a client application, the display treats the remote manager exactly as it treats the local one. If the user logs in, she receives the native desktop of the remote machine. The remote Display Manager interface generally resembles the local one, except that options applicable only to the local machine are not shown, and the user is provided with an option to return back to the login screen of her local machine.

If the user selects "Choose Host From List", the Display Manager will poll the administrative network with an XDMCP broadcast, to which servers running a remotely accessible xdm service may respond. The resulting list is displayed as a menu. When a remote host is chosen, the local display manager will use XDMCP to hand control over to the remote host as described above. The XDMCP probe broadcast is only addressed to the local data-layer (e.g. ethernet) subnetwork. A server that is accessed from the workstation through a router, or an IP (Layer 3) switch, can only be invoked through the "Enter Host Name" option rather than from "Choose Host From List".

When the user exits the remote desktop login, or chooses to return to the local login screen, control is transferred back to the local display manager, which re-displays the local login screen with its local options.

MANAGING UNIX/LINUX SERVERS FROM A MICROSOFT WINDOWS DESKTOP

Since the X Window System and XDMCP are publicly documented protocols, X Window System servers and XDMCP clients may be written for any platform, including most versions of Microsoft Windows. Several implementations of X servers and XDMCP clients for Microsoft Windows are widely used by server administrators who work from Microsoft Windows workstations. The most popular, at this time, is the free open-source implementation included in the Cygwin (<http://www.cygwin.com>) distribution. Cygwin is a volunteer open-source project sponsored by Red Hat, Inc., to permit remote administration of Red Hat Linux and other Linux or Unix servers from the Windows desktop.

The Cygwin distribution is based on the Cygwin Linux emulation layer for Microsoft Windows platforms. The Cygwin Linux emulation layer presents the interface of a DLL (Windows loadable library) to the Windows kernel, and the interface of a Unix kernel to applications originally designed to run in a Unix or Linux environment. The Cygwin distribution includes, in addition to the emulation layer, all the components of a standard Linux distribution, including an X server.

When used for remote server administration on a workstation with network interfaces to both the internal client network and an administrative network, the "-from local_host_name_or_ip_address_on_the_administrative_network" option must be specified when starting XWin.exe. A comprehensive section on XDMCP under Cygwin is included in the Cygwin FAQ, <http://x.cygwin.com/docs/faq/cygwin-x-faq.html>.

Since Unix/Linux components often depend on each other, and server administration workstations are likely to offer adequate hard disk space, it is a good idea for a server administrator to install the entire Cygwin distribution on her Windows workstation. This is most easily done by following the instructions in the "How do I just get everything?" section of <http://cygwin.com/faq.html>.

REMOTE DESKTOP PROTOCOL

On Microsoft Windows platforms, user interface functions are normally available to applications through calls to operating system libraries. The control of user interface hardware is an operating system function, and on traditional desktop versions of Microsoft Windows, application calls to user interface libraries resulted in direct execution of OS kernel code. In more recent versions of Microsoft Windows, the system library calls may be redirected to a remote user interface through the Remote Desktop Protocol (RDP) [Microsoft 2003] service. On Microsoft platforms, the application side is a service - that is, an implementation of server functionality - to which the desktop computer that provides the physical user interface connects, through the network, as a client. RDP is Microsoft's functional counterpart of the X Window System and XDMCP protocols combined.

Microsoft's RDP is nominally an extension of the ITU-T T.128 (or "T-Share") [ITU 1998b] application sharing protocol, using the ITU-T T.125 [ITU 1998a] multipoint communication service standard. RDP also uses parts of IETF RFC 905 [ISO DP 8073 Transport Protocol] [ISO 1984] and RFC 2126 (ISO transport over TCP) [Pouffary and Young 1997]. Unfortunately, no documentation of RDP itself is available publicly. In the absence of knowledge of the protocol, organizations outside Microsoft are not generally able to debug RDP configuration problems internally. The absence of complete public documentation also limits the ability of organizations outside Microsoft to evaluate the security of RDP in their networks. The latter problem is mitigated by isolating the organization's internal server administration network from other networks, as shown in Figure 4.

RDP clients are distributed, by Microsoft, for desktop platforms (such as Windows XP Professional) that may be used on server administration workstations. Administrators who use Unix or Linux based workstations may administer Windows based servers using the open source rdesktop package (<http://www.rdesktop.org/>); proprietary software such as CITRIX ICA (<http://www.citrix.com/site/PS/products/QA.asp?familyID=19&productID=1449&faqID=5638&featureID=QAP>); or by adding a second, Windows-based computer to their desktop workstation. Sun Microsystems desktop workstations, which are designed for server administration, accept an optional plug-in board that actually is an entire Microsoft Windows compatible computer. This plug-in second computer is configured to use the workstation's KVM hardware as its user interface. In the most common configuration, the native Windows desktop is displayed in a window on one of the display screens of the workstation. This window allows Microsoft servers in a mixed server network to be administered entirely with native Microsoft software, directly from a Sun Microsystems desktop workstation.

V. DISTRIBUTING THE ADMINISTRATOR'S FILESYSTEM

In server administration networks it is generally not necessary to replicate the administrator's home file system on individual servers. The administrator typically maintains a single set of files and directories on his or her workstation, and mounts this filesystem remotely, on each of the servers he or she administers, over the administrative network. The strict separation of the administrative network from other internal and external networks minimizes the risks and vulnerabilities (e.g., from share spoofing) that may result from mounting these filesystems remotely. A single home login filesystem, mounted where needed, gives the administrator uniform access to all her files and automatic configuration settings. Such functions as productivity optimizations and shell histories for command repetition will always be consistent and available regardless of which server the administrator is working on. Unix/Linux shells and window

managers can be configured so that the current hostname is always displayed in the window title bar, in every shell command prompt, and perhaps also through other indicators such as color and background schemes. These cues help the administrator stay aware of which server she is working on, while sharing a single personally customized working environment across servers.

Most information resources for server administrators, including sources of software examples and downloads are accessed via the World Wide Web. One of the advantages of a single remotely mounted login directory across all administered servers is that this provides a single, consistent bookmark ("favorites") file and other browser customization files. A reasonable security precaution is to let the administrator's browser access the Web only through a dedicated web proxy server for the server administration network. In this way all web resource access can be logged, and if necessary filtered for additional security.

The server administrator's primary workstation is usually a Unix or Linux system. The administrator's files are exported to other Unix and Linux servers using NFS [Sun Microsystems 1989], the Unix/Linux native Network File System sharing protocol. File systems from Unix/Linux machines can be exported to Microsoft Windows systems using Samba, an open-source Unix/Linux implementation of SMB (Server Message Block) [Sharpe 2002] and CIFS (Common Internet File System) [Microsoft 2004], the native Microsoft Windows file system sharing protocols.

If the server administrator's primary workstation runs a version of Microsoft Windows, the administrator's home filesystem will be shared with Windows-based servers using native SMB/CIFS over the server administration network. Sharing the administrator's home SMB/CIFS filesystem with the administrator's logins on Unix servers is usually done by setting up a dedicated Linux computer as a filesystem relay. The Linux kernel (but not most Unix kernels) supports Samba for mounting SMB/CIFS filesystems remotely. The remotely mounted SMB/CIFS filesystem is then exported to Unix-based servers with NFS. A small computer dedicated as a Linux filesystem relay box is usually adequate for distributing home filesystems from SMB/CIFS to NFS for all users of the administrative network.

To maximize productivity, it is advisable that not only server administrators but also webmasters, database administrators and other users of the server administration network arrange to have a single home filesystem on their desktop workstation, and export it to all machines they work on, using the server administration network.

VI. THE ADMINISTRATOR'S WORKSTATION

For optimum productivity, system administrators (and other IT personnel with network and server responsibilities, e.g. database administrators and webmasters) should use desktop machines designed for administrative work. While workstations are more expensive than ordinary PCs, the additional cost is quickly amortized by savings in loaded salary for administrator time. Multi-screen workstations designed for optimum productivity in server administration tasks were pioneered by Sun Microsystems, and the following description is based on the capabilities of their current workstation models such as Sun Blade 150. Over and above the capabilities of a typical personal computer, these workstations provide:

- Processor, operating system, and executable code formats identical with those for the vendor's servers, from the smallest single-processor server to the largest server cluster. Thus, software can be fully pre-tested on the administrator's workstation before it is installed on a production server from the same vendor.
- Support for multiple frame buffers to drive several display screens. The typical server administrator workstation uses three screens, usually with monitors arranged in an arc around the user's keyboard. These screens are positioned on furniture designed for this function. The keyboard and mouse surface are mounted on what furniture vendors call a

"corner unit" or "corner desk", which lets the user position the 3 monitors at 45-degree angles to each other. The keyboard surface is below desk-height to minimize the user's susceptibility to carpal-tunnel syndrome. The monitors are elevated above the surface of the desk to eye-level, and titled forward to reduce glare. The keyboard may provide extra columns of shortcut keys left of the alphabetic keys. An extra middle (or side) button on the mouse provides a one-press paste-at-cursor function.

- Software support for server administrator productivity. This support includes access to xdmcp functionality, for bringing the native user interfaces of remote servers to the administrator's desktop, in the standard login screen. The multi-screen X Window System server provides smooth movement of the mouse cursor across the several screens, and also provides a single cut-and-paste buffer which is shared across all screens and all application windows, including remote applications connected from remote machines. The X Window System display screens can also display Microsoft Windows applications running on an optional PC-on-a-card plug-in board; these applications also have access to files on the workstation's file system, and to the cut-and-paste buffer for transfer of information to and from other local and remote applications.

With three screens, the application windows of several applications can be arranged so that they are visible simultaneously, saving time that would otherwise be wasted switching among windows. In Unix/Linux environments additional time is saved by allowing keyboard focus to follow the mouse, so that it is not necessary to waste time clicking on windows to activate them. Selections are copied to the paste buffer automatically as they are made, and can be pasted to any application on any screen with the middle mouse button. A typical copy-and-paste operation that takes 6 steps on a Microsoft Windows desktop (click on donor window to activate it, select information, copy to edit buffer, click on receiving window to activate it, click to specify paste-in location, then paste) is performed on the server administrator's workstation in two (select and paste).

Figure 5. CDE 3-Screen Remote Administration of a Solaris Server

Figure 6. KDE 3-Screen Remote Administration of a Linux Server

Figures 5 and 6 show the administrator's working user interfaces on the same 3-screen workstation to two servers. Figure 5 shows the CDE interface of a Solaris 9 server. Figure 6 shows the KDE interface of a Red Hat 8 Linux server. To assure positive transfer of skill between different user interface environments, the layout and other customizations are kept as similar as possible. In both cases, the center screen is used to display two side-by-side shell command windows, also used for a file editor such as vi, a mail client such as mutt etc. If more than two shell windows are needed, several shell windows can be stacked in the same position on each side. Under CDE, the "Front" key on the workstation keyboard can be used to shuffle the stack to access the desired window. KDE includes an application, "konsole," designed for stacking and shuffling shell and other superposed application windows. The right-side screen displays the web browser, such as Mozilla. Again, several browser windows may be stacked, and shuffled – for example, in CDE with the "Front" key. The left-hand screen is used to display the server administrator's tools. In Figure 5, that screen is displaying the Solaris Management Console. In Figure 6, it shows the Red Hat Package Manager.

VII. SECURITY, OFF-SITE ACCESS, AND HUMAN IMPACT

Since remote administration uses services through which server computers could be compromised, it is essential that protocols used in remote administration (RDP, X and XDMCP) be blocked from crossing the boundary of the server administration network. As a further security precaution, access to X-related services (including X and XDMCP) through the administrative network interface should be restricted, with configuration files, to the network addresses of the administrators' workstations. If it is necessary to provide administrative access to the servers from outside the administrative network - for example, so that the server administrator can respond to off-hours emergencies from home - then the X Window System connections to the administrator's remote machine may be forwarded over an ssh (secure shell) connection, through an ssh boundary server (one of the servers in the Firewall and Boundary Servers group in Figures 2 through 4,) the internal client network, and the administrator's workstation. Current operating systems for administrator workstations are pre-configured with an ssh (Secure Shell) daemon for this purpose. Ssh can present to the server administration network the interface of an X Window System server running on the administrator's workstation, but forward the actual X Window System traffic – and the complete user interface to administrative applications running on the administered servers, the SNMP management client, the web browser running on the administrator's workstation etc – over a securely encrypted tunnel to the administrator's home or other off-site (e.g. emergency) location.

When the server administration network architecture described in this article is applied so that each service administrator interacts with servers only through the physical user interface of a single workstation, the server machines may be located in any appropriate location, such as a telecom's closet or a secure, climate controlled room in the service core of a business building. As a result, the location of the server machines can be optimized for physical security without the constraint of needing to provide a tolerable occupational environment to one or more co-located humans. There is no problem even if the machines are located in a place that is noisy, lacks ventilation, or is held at a temperature that would not be tolerated by employees. In case of extreme security needs the entire server farm may be located remotely, and can be physically secured against earthquakes, terrorist actions etc. User interface delivery through a small number of administrative workstations also permits logging of administrator activity to any extent necessary, all the way to a complete audit trail, for example logging of X Window System activity including logging of all keyboard, display and mouse events. There are business contexts in which this kind of complete audit trail will help minimize security risks and legal exposure.

Additional advantages of the server administration network architecture are its impact on the productivity, motivation and morale of server administration personnel. A single highly optimized user interface will minimize task time. Delivering this user interface through a single desktop workstation eliminates time for such tasks as walking to server machine locations or between machines. A multi-screen administrative workstation, and the autonomy to configure and

customize it for one's own maximum productivity and comfort, conveys institutional recognition of the server administrator's competence and importance to the business.

VIII. IMPLICATIONS FOR IS CURRICULA

Nearly all careers, for which students in Information Systems curricula are being prepared, are in some degree affected by network partitioning and by the deployment of dedicated server administration networks. System, network, and security administrators need the knowledge presented in this tutorial both to understand what they are doing, and to work effectively and productively with the networks they will be managing and using. Database administrators and software developers will need to understand the security function of network partitioning, to make sure that every action with security implications is authenticated not only with respect to the identity and role of the user or software agent performing it, but also with respect to the identity of the originating network interface, and of the peer host of the corresponding network connection. System planners and integrators, and organizational IT managers, will need this information to take into account the network context of performance and security, where nearly all contemporary business and organizational software is network-dependent or network-based, and network-administered.

For these reasons Figure 4 of this tutorial, or its equivalent, needs to be presented and explained at the beginning of every IS course and course segment that deals with network, server administration, or security issues. Courses dealing with server, network, or security administration ought to discuss essentially everything in this tutorial at some point in the curriculum. Familiarity with at least the equivalent of Figure 4 should be expected, and included, in all degree and certification examinations in Information Systems and Information Systems Management.

Information Systems careers that use the Server Administration Network will become increasingly desirable in the coming years. Of all technical careers in Information Systems, server, network and security administration jobs are least amenable to export overseas. Unlike user desktop applications, which are designed for ease of learning and support, and are only now starting to become automated, server, network and security administration were designed from the start to be as automated as possible – so that loss of jobs to automation is likely to be minimal. These factors are likely to boost enrollment in courses relevant to these specialties. My hope is that the present tutorial will help IS faculties to meet this challenge.

Editor's Note: This article was received on September 2, 2004. It was with the author for 5 revisions and was published on May 8, 2005.

REFERENCES

Editor's Note: The following reference list contains hyperlinks to World Wide Web pages. Readers who have the ability to access the Web directly from their word processor or are reading the paper on the Web, can gain direct access to these linked references. Readers are warned, however, that

1. these links existed as of the date of publication but are not guaranteed to be working thereafter.
2. the contents of Web pages may change over time. Where version information is provided in the References, different versions may not contain the information or the conclusions referenced.
3. the authors of the Web pages, not AIS, are responsible for the accuracy of their content.
4. the author of this article, not AIS, is responsible for the accuracy of the URL and version information.

- Braca, M. (1991) "X Display Management" *UNIX World* (8)4, pp. 107-112, April
- Case, J., M. Fedor, M. Schoffstall, and J. Davin (1990) "A Simple Network Management Protocol (SNMP)" IETF RFC 1157, May <http://www.ietf.org/rfc/rfc1157.txt> (current July 26, 2004).
- Cheswick, W. R. and S. M. Bellovin (1994) *Firewalls and Internet Security: Repelling the Wily Hacker*. Reading, MA: Addison-Wesley
- Egevang, K. and P. Francis (1994) "The IP Network Address Translator (NAT)" IETF RFC 1631, Cray Communications - NTT, May 1994 <http://www.ietf.org/rfc/rfc1631.txt> (current July 26, 2004).
- Fielding, R. et al. (1999) "Hypertext Transfer Protocol – HTTP/1.1" IETF RFC 2616, June 1999 <http://www.ietf.org/rfc/rfc2616.txt> (current July 26, 2004).
- ISO (International Standards Organization) (1984) "ISO Transport Protocol Specification, ISO DP 8073" IETF RFC 905, April <http://www.ietf.org/rfc/rfc905.txt> (current July 26, 2004).
- ITU (International Communications Union) (1998a) *ITU-T Recommendation T.125, Multipoint Communication Service Protocol Specification*. Geneva: International Communications Union, February.
- ITU (International Communications Union) (1998b) *ITU-T Recommendation T.128, Multipoint Application Sharing (T-Share)*. Geneva: International Communications Union, February.
- The KDE Team (2002) "K Desktop Environment, Revision 3.00.00" <http://www.kde.org/documentation/userguide/> (current July 26, 2004).
- Microsoft (2003) "Remote Desktop Protocol Specification Now Available Under License" <http://www.microsoft.com/windows2000/server/evaluation/news/bulletins/rdpsspec.asp> (current July 26, 2004).
- Microsoft (2004) "Windows 2000 Network Architecture: Common Internet File System" http://www.microsoft.com/windows2000/techinfo/reskit/en-us/cnet/cnad_arc_endh.asp (current July 26, 2004).
- The Open Group (1998) "User Interface: Common Desktop Environment" Document Number X98XD, <http://www.opengroup.org/bookstore/catalog/x98xd.htm> (current July 26, 2004).
- Packard, K. (2004) "X Display Manager Control Protocol, Version 1.1, X.Org Standard, X Version 11, Release 6.7", March 11 <http://www.x.org/xdmcp.ms> (current July 26, 2004).
- Postel, J. B. (1982) "Simple Mail Transfer Protocol" IETF RFC 821, August <http://www.ietf.org/rfc/rfc0821.txt> (current Jul. 26, 2004).
- Pouffary, Y. and A. Young (1997) "ISO Transport Service On Top Of TCP (ITOT)" IETF RFC 2126, March <http://www.ietf.org/rfc/rfc2126.txt> (current July 26, 2004).
- Rescorla, E. (2000) "HTTP over TLS" IETF RFC 2818, May <http://www.ietf.org/rfc/rfc2818.txt> (current July 26, 2004).
- Scheifler, R. W. and J. Gettys (1992) *X Window System: The Complete Reference to Xlib, X Protocol, ICCCM, XLFD. Third Edition: X Version 11, Release 5*. Burlington, MA: Digital Press.
- Sharpe, R. (2002) "Just What Is SMB? V1.2" <http://www.samba.org/cifs/docs/what-is-smb.html> (current July 26, 2004).
- Smith, R. W. (2002) *Advanced Linux Networking*. Boston, MA: Addison-Wesley.
- Sommerfeld, W. (2003) "Secure Shell (secsh)" <http://www.ietf.org/html.charters/secsh-charter.html> (current July 26, 2004).
- Sun Microsystems (1989) "NFS: Network File System Specification" IETF RFC 1094, March 1989 <http://www.ietf.org/rfc/rfc1094.txt> (current July 26, 2004).
- Warkus, M. (2004) *The Official GNOME 2 Developer's Guide*. San Francisco, CA: No Starch Press, April.

Wilder, R. et al (2004) "Layer 2 Virtual Private Networks (l2vpn)" <http://www.ietf.org/html.charters/l2vpn-charter.html> (current July 26, 2004).

GLOSSARY

Artisanal

Craftsmanly; artisanal knowledge is knowledge that is conveyed by supervised practice and collegial collaboration, as is done among craftsmen, rather than by formal documentation or instruction.

Boundary Gateway

See "Boundary Router."

Boundary Router

(Also "Gateway Router" or "Boundary Gateway.") The router between an organization's networks and one or more networks external to the organization, such as the Internet.

Buffer Overflow

Software defect that permits a program to change the content of memory beyond the area allocated for the variable whose value the program is supposed to be changing. This defect may permit an attacker to change the values of other variables, leading to consequences not intended by the programmer. For example, an attacker may be able to change the string to which an authorization password will be compared, allowing the attacker to obtain unauthorized privileges on the vulnerable system.

CDE

Common Desktop Environment. A software package consisting of a window manager, user interface libraries, and customization files that emulated the largely Microsoft-compatible user interface of IBM's OS-2 under the X Window System. Until Gnome and KDE became available for server-based workstations, CDE was often used by server administrators to maximize positive transfer of skills and habits between Microsoft and Unix platforms.

CIFS

Common Internet File System. A proprietary Microsoft filesystem designed to permit sharing over a network with the Server Message Block protocol. See "SMB."

Co-Location

The placement of computers and personnel in close physical proximity, to permit high-speed networking and physical access to computer hardware.

Console-Based Administration

Computer system administration using the computer's built-in user interface hardware - usually built-in interfaces for a keyboard, a video monitor, and a mouse - "KVM."

Demilitarized Zone

The network segment between the external firewall and outward-facing organizational servers.

Display Manager

An X Window System client that controls an X Server between login sessions, and provides a user interface for remote graphical login.

DMZ

See "Demilitarized Zone."

Filesystem

A hierarchical system of directories (folders) containing computer files, under a single top-level directory, whose name is used as the name of the filesystem.

Gateway Router

See "Boundary Router."

Gnome

GNU Network Object Model Environment. A free, open-source software package consisting of a window manager, user interface libraries, and customization files that provide a user interface based on Microsoft-compatible user habits and skills under the X Window System.

Graphical Shell

A graphical user interface to a computer operating system. Microsoft operating systems provide a built-in graphical shell called a "Desktop." On Unix and Linux systems, the Window Manager also serves as a Graphical Shell.

HTTP

Hypertext Transfer Protocol. A protocol and service that provides files and processing resources to clients such as World Wide Web browsers. HTTP is the main session-layer protocol of the World Wide Web.

HTTPS

Hypertext Transfer Protocol over Secure Socket Layer. HTTP secured by a cryptographic protocol that provides greater privacy and improved authentication. HTTPS is the main session-layer protocol of electronic commerce and automated web services.

IP

Internet Protocol. An addressing and delivery protocol that enables worldwide host-to-host (computer to computer) communication through the interconnected networks that form the global Internet.

KDE

The K Desktop Environment. Originally an effort to develop a desktop environment compatible with proprietary CDE, but for open-source Linux and BSD Unix operating systems, KDE is now a widely used alternative to Gnome. KDE and Gnome use similar user interface conventions, so that KDE applications are readily usable on Gnome and vice-versa.

KVM Switch

A device designed to switch a single keyboard, video monitor, and mouse ("KVM") among the keyboard, video monitor and mouse interfaces of several computer systems.

Mail Client

An application that provides a user interface for reading and sending electronic mail.

NAT

See "Network Address Translation."

Network Address Translation

A scheme to hide internal network addressing and traffic behind a single externally visible IP address, so that the combined external traffic of one's whole organizational network, or a large segment of one, looks to the external Internet like traffic from the network interface of a single host computer. With NAT, internal machines that do not need to receive incoming external connections use IP addresses drawn from a separate, private IP address space (usually a subset of 192.168.*.*) The vectors of IP addresses and TCP port numbers of outbound connections from those (possibly hundreds of) internal machines are translated to TCP port numbers on the single externally visible IP address.

Network Interface

The connector and circuit that connects a computer to a specific data network.

NFS

Network Filesystem. An open protocol traditionally used to share filesystems among Unix and Linux computers. Also optionally available under some Microsoft server operating systems.

Packet Filter Firewall

An early firewall design that filtered traffic according to "permit" and "deny" rules based on packet attributes.

RAM

Random Access Memory. Information storage circuits designed so that any information contained in them can be accessed with the same, very short delay. Software processes normally interact only with information that is present in RAM; information that is not already present must be brought into RAM from its location in slower but more permanent storage, such as a hard disk, before it can be used.

Remote Desktop Protocol.

A proprietary Microsoft protocol that permits a Microsoft operating system desktop to be displayed on, and controlled from, a remote host computer.

Shell

An application that provides a user interface to a computer operating system. See "Graphical Shell" and "Shell Programming Language."

Shell History

A file that stores a history of shell commands. See "Shell Programming Language."

Shell Programming Language

Popular Unix and Linux command-line shells [Bourne Shell, C Shell, Korn Shell, Bourne-Again Shell etc) support command languages each of which is also a programming language. Commands are stored in a Shell History file, so that any sequence of previously issued commands can be edited into a shell program, often called a "shell script," that can be executed automatically (even without human intervention) like any other computer program.

SMB

Server Message Block. A proprietary Microsoft protocol for sharing filesystems and print services among computers on a network.

SMTP

Simple Mail Transfer Protocol. A protocol and service that accepts electronic mail messages from clients (or intermediate SMTP servers) and forwards electronic mail to its destination.

SNMP

Simple Network Management Protocol. The network administrator's SNMP-based network management software - the "SNMP Manager" - collects information from "SNMP Agents" built into each "managed" network device, such as a switch, a router etc.

SSH

Secure Shell. A protocol and service that provides remote command line login and execution, file transfer, and forwarding of otherwise insecure protocols over the Secure Socket Layer for greater privacy and improved authentication.

Stateful Firewall

A firewall that monitors the establishment of authorized TCP connections, and stops inbound packets containing TCP segments that do not belong to previously established authorized TCP connections.

Tunnel

A forwarding arrangement that authenticates and encrypts messages of otherwise vulnerable protocols for more secure delivery. Tunnels are used to permit reasonably secure continued use of network software that was developed before the emergence of security concerns.

VPN

Virtual Private Network. A protocol that provides Tunnels (which see) for secure transmission of communications that otherwise would have been vulnerable to interception or forgery; a network that uses a tunneling protocol for this purpose.

Window Manager

An X Window System client application that provides a user with a graphical interface to the capabilities of an X Server, and of the underlying operating system.

XDM

X Display Management. See "Display Manager."

XDMCP

X Display Manager Control Protocol. The X Window System protocol that provides remote graphical login capabilities. See "Display Manager."

X Server

Software that provides X Window System services to local or networked clients. See "X Window System."

X Window System

A platform-independent protocol and service that provides its clients (X Window System applications) with access to the user interface hardware (a keyboard, a pointer device such as a mouse or trackball, and one or more display screens) of a computer or terminal. An X Window System server - "X Server" for short - is a standard component of every Unix or Linux operating system, and may be added to most other operating systems.

XWin.exe

The conventional name of the X Server process on Microsoft platforms. See "X Server."

ABOUT THE AUTHOR

Adam Reed is Associate Professor of Information Systems at California State University, Los Angeles. His current research interests include productivity and accessibility of information systems, epistemology of knowledge representation, programming languages, operating systems, network protocols, and network security administration. His research publications have appeared in *Science, Memory and Cognition, Contemporary Psychology, Journal of Verbal Learning and Verbal Behavior, Behavior Research Methods and Instrumentation*, and in *Communications of AIS*. He is the author of three chapters on Unix in Rosen et al. [1996], He holds 3 US Patents.

Copyright © 2005 by the Association for Information Systems. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, P.O. Box 2712 Atlanta, GA, 30301-2712 Attn: Reprints or via e-mail from ais@aisnet.org.

Communications of the Association for Information Systems

ISSN: 1529-3181

EDITOR-IN-CHIEF

Paul Gray

Claremont Graduate University

AIS SENIOR EDITORIAL BOARD

Detmar Straub Vice President Publications Georgia State University	Paul Gray Editor, CAIS Claremont Graduate University	Sirkka Jarvenpaa Editor, JAIS University of Texas at Austin
Edward A. Stohr Editor-at-Large Stevens Inst. of Technology	Blake Ives Editor, Electronic Publications University of Houston	Reagan Ramsower Editor, ISWorld Net Baylor University

CAIS ADVISORY BOARD

Gordon Davis University of Minnesota	Ken Kraemer Univ. of Calif. at Irvine	M.Lynne Markus Bentley College	Richard Mason Southern Methodist Univ.
Jay Nunamaker University of Arizona	Henk Sol Delft University	Ralph Sprague University of Hawaii	Hugh J. Watson University of Georgia

CAIS SENIOR EDITORS

Steve Alter U. of San Francisco	Chris Holland Manchester Bus. School	Jaak Jurison Fordham University	Jerry Luftman Stevens Inst. of Technology
------------------------------------	---	------------------------------------	--

CAIS EDITORIAL BOARD

Tung Bui University of Hawaii	Fred Davis U. of Arkansas, Fayetteville	Candace Deans University of Richmond	Donna Dufner U. of Nebraska -Omaha
Omar El Sawy Univ. of Southern Calif.	Ali Farhoomand University of Hong Kong	Jane Fedorowicz Bentley College	Brent Gallupe Queens University
Robert L. Glass Computing Trends	Sy Goodman Ga. Inst. of Technology	Joze Gricar University of Maribor	Ake Gronlund University of Umea,
Ruth Guthrie California State Univ.	Alan Hevner Univ. of South Florida	Juhani Iivari Univ. of Oulu	Claudia Loebbecke University of Cologne
Michel Kalika U. of Paris Dauphine	Munir Mandviwalla Temple University	Sal March Vanderbilt University	Don McCubbrey University of Denver
Michael Myers University of Auckland	Seev Neumann Tel Aviv University	Dan Power University of No. Iowa	Ram Ramesh SUNY-Buffalo
Kelley Rainer Auburn University	Paul Tallon Boston College	Thompson Teo Natl. U. of Singapore	Doug Vogel City Univ. of Hong Kong
Rolf Wigand U. of Arkansas, Little Rock	Upkar Varshney Georgia State Univ.	Vance Wilson U. of Wisconsin, Milwaukee	Peter Wolcott U. of Nebraska-Omaha
Ping Zhang Syracuse University			

DEPARTMENTS

Global Diffusion of the Internet. Editors: Peter Wolcott and Sy Goodman	Information Technology and Systems. Editors: Alan Hevner and Sal March
Papers in French Editor: Michel Kalika	Information Systems and Healthcare Editor: Vance Wilson

ADMINISTRATIVE PERSONNEL

Eph McLean AIS, Executive Director Georgia State University	Reagan Ramsower Publisher, CAIS Baylor University
---	---