LANTRONIX®

KVM over IP for the "Distributed IT" Environment

Flexible, Scalable, and Cost-Effective Management for Remote Servers and Equipment

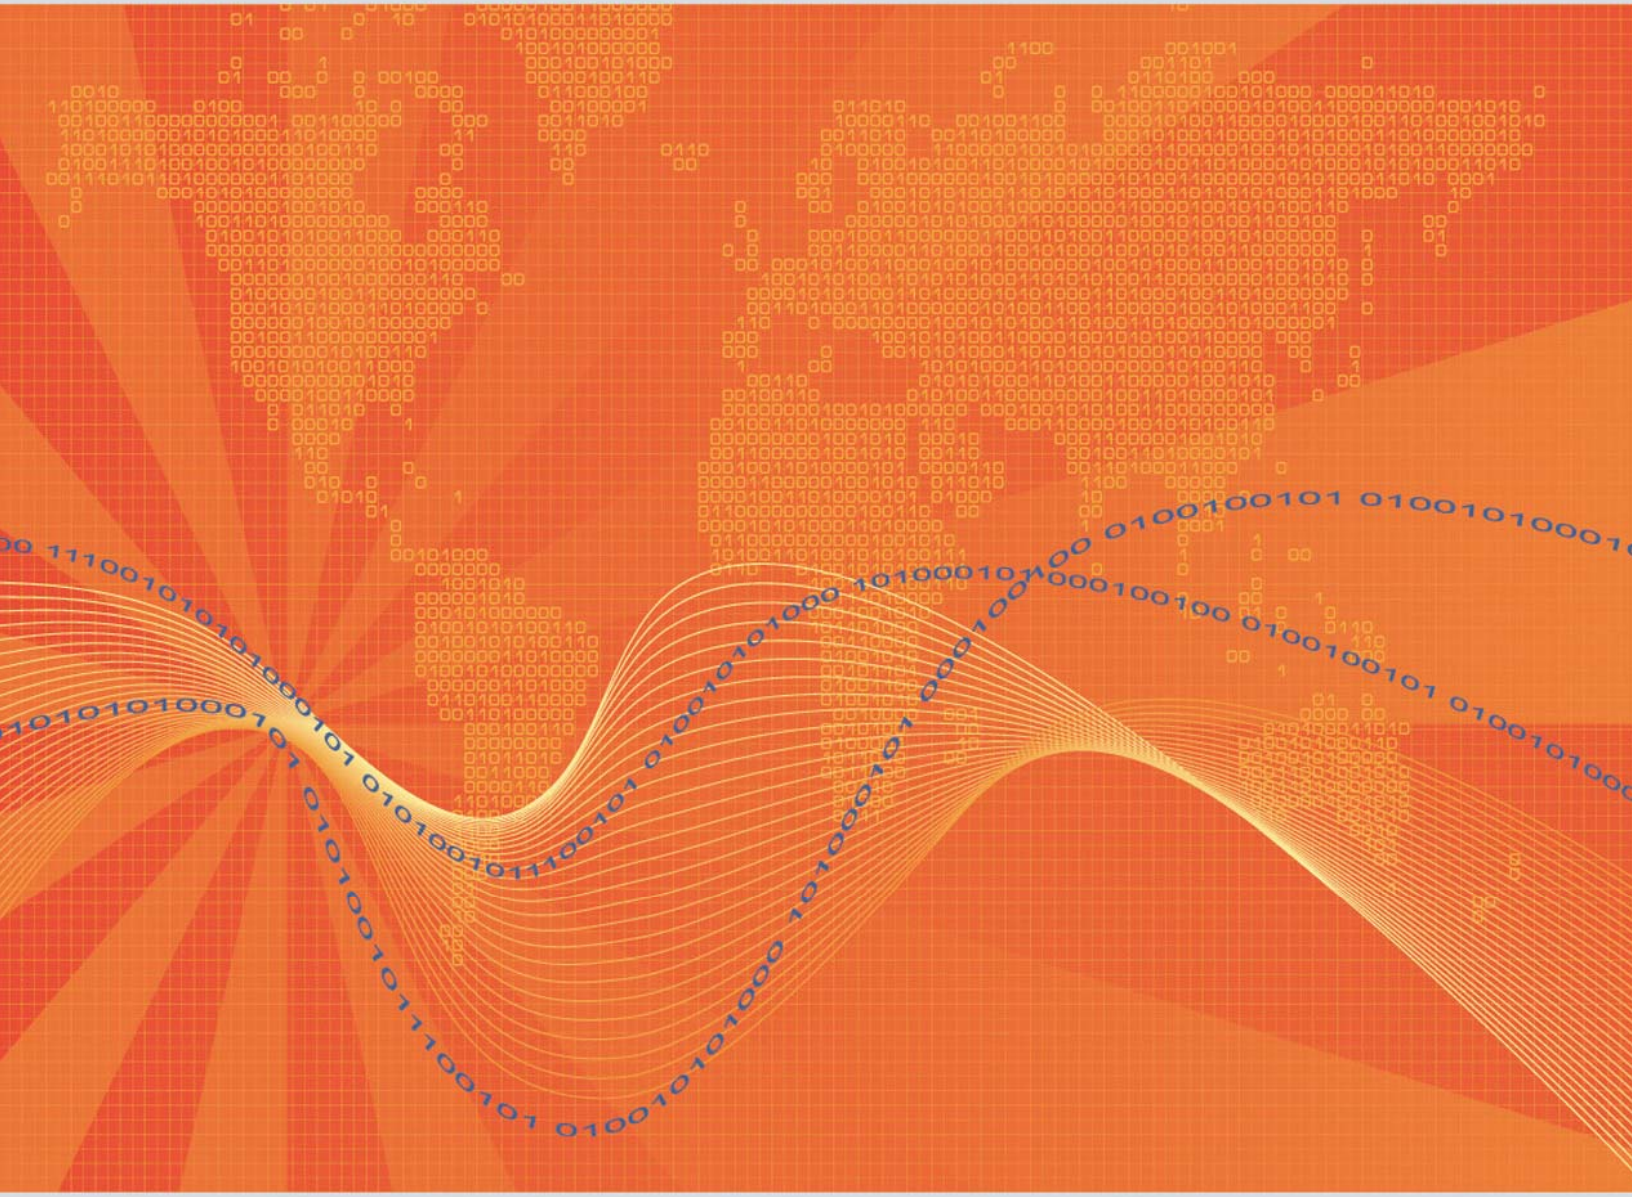

Lantronix, Inc.
15353 Barranca Parkway
Irvine, CA 92618
Tel: +1 (800) 422-7055
Fax: +1 (949)450-7232
www.lantronix.com

Contents

Introduction

With 27% of all businesses and an estimated 97% of large enterprises having branch offices making up some six million satellite facilities[1], the demand on IT departments supporting geographically dispersed networking infrastructure continues to grow. And it's not just branch locations that are fueling the demand; the very nature of business today demands a level of constant connectivity never seen before. In order to function effectively and stay competitive in the global arena, people at branch offices, remote facilities – or just on the other side of the corporate campus – require access to the same information and systems available at main headquarters. This demand creates a dispersed network environment with widely distributed IT assets, which in turn creates a greater administrative challenge for maintaining equipment to ensure network uptime. To meet this demand, organizations take advantage of remote infrastructure management equipment to access and manage servers and IT equipment over the Internet.

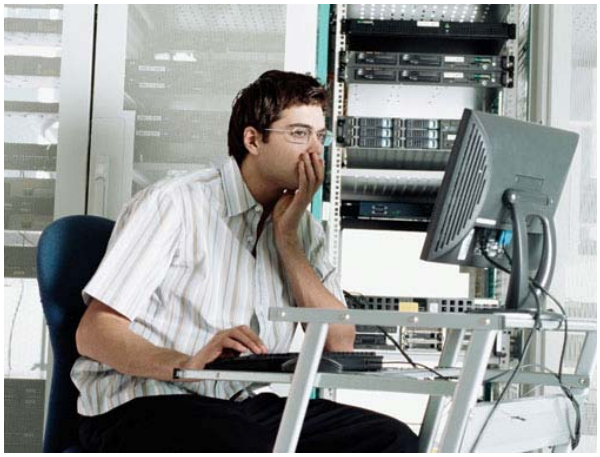

KVM (keyboard/video/mouse) over IP is an effective tool for remotely managing servers regardless of their physical location. With BIOS-level access, system administrators can monitor and respond to server issues from virtually anywhere over an Internet connection. For this reason, server management with remote KVM switches have become the cornerstone of most data centers today.

However, for all its benefits, KVM has a number of limitations when placed in the "distributed IT environment" where equipment is dispersed throughout a large building, across a campus or even across the globe. These limitations include:

- *Limited configuration flexibility and expandability:* Traditional KVM over IP switches come in multi-port configurations (8/16/32, for example). In the branch office environment however, there may only be two servers, leaving some unused ports and increasing the cost-per-port. Or, adding an additional server to an even port count can force administrators to purchase an extra multi-port switch when only a single additional port is required.
- *Potential blocked access to critical servers:* Typical high-density KVM over IP solutions allow access to a limited number of managed servers at a time. For example a high-end 8-port KVM may only allow one or two servers to be accessed at a time. In a busy IT environment, this could prevent an administrator from performing needed maintenance or service.
- *Reliability and access to multiple servers:* If access to all 32 servers is through a single Ethernet connection and the switch port fails, access to all 32 servers is lost.
- *Need for extra hardware and software:* Switch-based KVM over IP solutions may require separate KVM and serial "dongles" to be attached to the server(s). Special

[1] [1] Source: IDC

client-based software may also be required. And, additional external power supply(s) may be necessary.

- *Distance limitation between servers and the KVM switch:* KVM over IP switches have a CAT5 cable length distance limitation of approximately 50 to 150 feet between the server and the KVM before the analog signal begins to degrade, limiting the flexibility in a distributed IT environment.

Because of these considerations, IT professionals should understand the challenges associated with implementing KVM, as well as the differences between "high-density" and "distributed KVM" in order to select the best technology solution for their environment. This paper will present the complexities of managing the distributed data center and highlight the advantages of using a KVM over IP solution that delivers flexible, scalable and affordable CAT5-based remote access.

IT Changes, IT Stays the Same

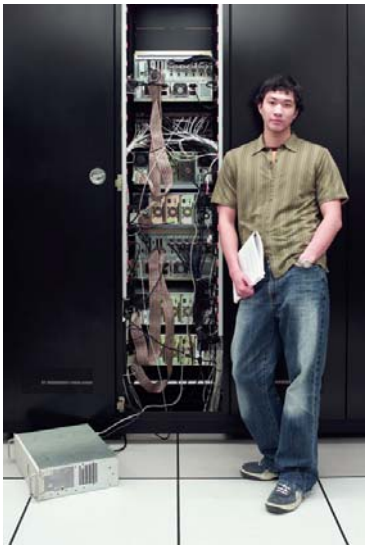

Today's IT professionals continue to be challenged by unrelenting changes in the enterprise: explosive data growth, more compliance regulations, increased application complexity, geographic distribution of assets, and expanding SLAs[2] with tighter MTTR[3] requirements, just to name a few. Business continuity is *the critical requirement* for companies of all sizes to remain competitive, with 24/7 application uptime and secure global 'anytime anywhere' information access, which is a fundamental operational need that has fallen on the IT team to design, support, maintain and deliver. The financial and operational repercussions of system downtime continue to be very severe, resulting in reduced employee productivity, regulatory penalties, added costs and large revenue impacts. A 2006 study estimated the combined cost of one hour of downtime cost the average business more than $1 million USD[4].

Additionally, a growing number of business applications are now mission critical, consuming additional IT personnel and budget resources that have not kept pace with changes. Take company e-mail as an example. As much as 75% of a company's intellectual property is contained in e-mail messages[5], with IT administrators spending more than 25% of their time managing their e-mail infrastructure alone[6].

With all these changes, some things have remained unchanged. IT departments are still expected to develop strategic and tactical plans that control costs, mitigate operational risk and improve data protection and information availability.

[2] Service Level Agreements
[3] Mean Time To Recovery
[4] META Group, Gartner 2006
[5] Source: Enterprise Strategy Group (ESG)
[6] Source: Coalition for Networked Information (CNI)

Whether an organization is large or small, maintaining and managing a secure and continuous distributed information infrastructure with limited resources is a challenge every IT professional faces today.

The (Re) Distributed Data Center

The fundamental concept of the data center is changing rapidly and dramatically. Primary data center elements including applications, servers, infrastructure and storage are being virtualized and redistributed in order to lower cost and complexity, improve asset utilization, and operational efficiencies.

Hosted applications, software as a service (SaaS), server-based appliances, grids and utility computing are becoming ubiquitous. Small and medium companies now routinely utilize server virtualization and network storage technologies, large enterprises are rapidly consolidating Tier 1 assets and operations, while growing branch office and remote production facilities. A recent article in a prominent IT publication stated that more than half of a company's data and servers reside *outside* what we consider to be the data center[7]. The sheer number of remote branch offices (ROBOs) with servers and applications requiring IT management is astounding: over six million branches and over 8 million small businesses, with 27% of all businesses having branch offices, and an estimated 97% of large enterprises having branch office facilities[8].

In addition to ROBOs and large scale geographic distribution of IT assets requiring 24/7 management, small scale campus and building IT infrastructure deployments are becoming increasingly interconnected, and have identical uptime and information availability requirements as the Tier 1 enterprise and ROBOs. According to a recent study[9], educational institutions are increasing spending on internetworking solutions to take advantage of the cost benefits offered by sharing resources amongst geographically distributed institutions. This has been cited as the number one top trend in higher education for networking between various branches and locations using LAN, WAN, external access, and the Internet. Meanwhile, IT professionals are spending more time managing security and other day-to-day challenges.

---

[7] Network Computing, 2006
[8] Source: IDC
[9] Frost and Sullivan, 2006

Even as IT assets become more distributed and interconnected, it is impractical for IT professionals to be similarly redistributed. It is still more cost effective to maintain centralized teams of professionals to provide IT support and management. The result is a lack of IT staff on-site and 'at the rack' for remote and branch office locations. With an absolute requirement to maintain system and application uptime everywhere, a primary administrative challenge continues to be how to cost effectively deploy, manage, maintain and troubleshoot geographically distributed servers and their applications with limited and centralized IT personnel.

Servers Everywhere

Server management is the foundation of any IT management strategy. If a server fails, applications fail, and operations come to an immediate halt. As the data center becomes virtualized and redistributed, the purposes, location and form factor of servers also changes, and the job of server management continues to evolve. In the Tier 1 enterprise, server density is increasing both on the hardware side with smaller and more compact form factors such as blade servers, and on the software side with growing adoption of server virtualization technologies such as VMware™. In both enterprise and ROBO environments, the use of servers as a platform has evolved from the days of core and layered production servers to a broad array of data center server functions that include:

- E-mail / 'Post office' servers
- List servers
- FTP Servers
- Certificate servers
- Backup servers
- Content and information management servers
- Departmental servers
- Proxy servers
- Database servers
- DHCP/WINS servers
- DNS (lookup) servers
- Rich Media servers
- SMTP servers
- HTTP/Web servers / hosts
- Redundant servers
- VPN and Gateway servers
- Java hosting servers
- Application servers
- Communication servers
- Fax servers
- File servers
- News servers (e.g. Usenet)
- Standalone servers
- Specialty Appliances (server based)
- Test and Development servers
- Dedicated data processing machines
- Control server

In addition to the growth in the functions and types of servers, the location of servers is expanding beyond the consolidated data center with racks and racks of servers:

The challenge of cost-effectively managing geographically dispersed servers has been addressed in part by the introduction of KVM over IP solutions that provide secure remote keyboard, video, mouse control of servers over an IP network. A KVM switch allows a single keyboard, video display monitor, and mouse to be switched to any of a number of computers at once, as opposed to having a single person interact with a number of computers one at a time. KVM switches are commonly used at Web and other server locations with multiple computers but can be managed with usually a single administrator or Webmaster. The switch provides more table space in addition to saving the cost of multiple keyboards and monitors.

Servers Everywhere	
Campus environments	Factories
Mid-size businesses	Government facilities
Distributed IT infrastructures (multi-floor, multi-building, multi-site)	Educational facilities
Distributed facilities	Dark data center facilities
Remote sites	Convention centers
Remote branch offices (ROBOs)	Corporate offices
Departmental IT	Kiosks
Labs: computer test labs, call centers, help desks, training rooms	Small and medium business (SMB)
Test and development (engineering IT)	Disaster recovery sites

Server Management in the Distributed Data Center

In the distributed data center, there are many elements to consider when developing an overall server IP-based KVM management strategy.

Guaranteed Anytime Anywhere Access. A server management solution should support guaranteed access (anytime, anywhere) to servers regardless of how many users are logged in. Criteria to look for include:

- Secure administrator access to servers over an IP network - from LAN or the Internet
- Non-blocked access to servers (e.g. no conflict for remote access)
- BIOS level access to servers
- No limitations on locations of servers (across the room, across the globe)

Simple and Non-Invasive. Any server management strategy should reduce complexity and lower the overall IT workload. Deployment, installation and usage of a server management solution should minimize or eliminate impact on the server, its operating systems, services and applications. The solution should offer:

- Simple and quick installation
- No special hardware (management cards, etc.) to install in server
- No special software (drivers, application software) to install and maintain on server
- Operating system and application independent
- Simplified remote maintenance (e.g. Virtual Media support)
- No special software to install and maintain on clients
- Small size (form factor)
- Minimal power consumption and heat generation
- Flexible cabling (utilize CAT5 cabling)
- No special training required to install or operate

Flexible, Scalable and Cost Effective. The dynamic nature of the distributed data center and cost constraints demands any remote server management solution to provide flexibility and scalability as part of the solution, without requiring the purchase of additional elements. A distributed solution should provide:

- No special software licenses to buy or maintain (client software, remote users, etc.)
- No special cabling limitations (supports CAT5 cabling)

- No special hardware or dongles to purchase and install
- Incremental 'add-as-you-grow' capability with no sudden cost increase to scale and add additional servers or remote administrators
- No conflict with existing server management solutions (no 'rip-and-replace')
- Ability to enhance existing server management solutions (local KVM)
- Ability to integrate into larger remote (OOBI[10]) management architecture

Fault Tolerant and Secure. Distributed remote assets must be fault-tolerant and secure by providing the following benefits:

- Allows management access when server, OS, or network stack is compromised
- No single point of failure (e.g. unit failure doesn't effect other systems)
- Designed for high MTBF[11] (e.g. no moving parts, cooling fans, power supplies)
- No special security software or infrastructure to purchase, install and maintain
- Uses existing IT security services (RADIUS, LDAP, Active Directory)
- Optional out-of-band access in case of network failure (e.g. serial modem access)

High-Density KVM or Distributed KVM

Traditional data centers with racks and racks of servers have primarily used high-density KVM switch solutions which take up one to two units of rack space and allow a connection to many servers. These solutions are tried and true and will continue to be used in situations where high-density servers reside.

However, these systems usually require a capital investment for the hardware and often require dedicated management software which, in some cases, also requires dedicated servers to run the software even to the extent that multiple software and server licenses are required for systems on different subnets. These solutions do not offer the flexibility and scalability needed for the changing landscape of distributed infrastructures.

The trend toward distributed environments with many servers deployed throughout the enterprise on different floors and extending outwards to remote branches requires a new way of thinking. Creative solutions for leveraging the Internet to access servers without the need to purchase high-density appliances or surrender to expensive software license contracts are now a requirement.

Determining the right solution greatly depends on the type of environment, the direction of the IT organization, and the underlying need for flexibility and scalability. For many organizations, a co-existence of high-density and distributed KVM will be the solution where flexibility and scalability are needed to augment existing high-density systems. For the growing trend of distributed computing and organizations with larger remote branch and campus environments, this new paradigm of distributed KVM with all its benefits will take the lead.

[10] Out Of Band Infrastructure
[11] Mean Time Between Failure

Distributed KVM for the Distributed Data Center

Lantronix SecureLinx Spider™ provides secure KVM (keyboard, video, mouse) management of servers over an IP network. Unlike traditional KVM switches, Spider offers a flexible, scalable and affordable CAT5-based remote access KVM solution in a cable-friendly, compact "zero-footprint" package.

The latest addition to the SecureLinx family of Device Management and Control solutions, this KVM over IP solution eliminates server-to-switch CAT5 cable distance limitations, and gives system administrators non-intrusive and cost-effective 24/7 access to servers across a wide variety of IT/network environments: mission critical servers in high-density data centers, servers distributed over corporate campuses, multi-floor buildings, remote/branch office sites, government facilities, convention centers, POS/kiosks, etc. Management access, from BIOS to applications, from any web browser anywhere, at any time is guaranteed.

Amazingly scalable, Spider is a complete KVM switch that can be easily daisy-chained together using Lantronix SwitchPort+™ integrated Ethernet switch technology. This provides a cost-effective solution in environments where numerous cable drops and distance limitations can be a challenge when adding servers. Rather than adding another more costly switch that forces administrators into pre-defined multiples, administrators can add remote IP users one server port at a time by simply deploying another Spider. And network reliability increases as there is no single point of failure with Spider deployments. If a unit is disabled, only access to that server is lost until the Spider is replaced.

SecureLinx Spider provides continuous availability to servers with 1:1 non-blocked, BIOS-level access. This allows administrators to have guaranteed access to mission critical servers regardless of how many of them need remote access. In other words, administrators are not "locked in" to a fixed number of remote users, and Spider offers an extremely low-cost-per-remote-user for guaranteed non-blocked access. And no client software or external power supply is required.

Benefits

- Full non-blocked access providing one of the lowest cost-per-remote-user server management solutions available
- Secure, full BIOS-level control of servers over an IP network
- Clean, flexible CAT5 KVM solution with virtually no cable length restrictions
- Completely integrated IP-based KVM switch in a "Zero U" form factor
- Cost-effective and easy "add-as-you-grow" scalability by daisy-chaining multiple units
- Compact server-powered design – no external power supply required
- Virtual Media support
- Browser-based – no client software required

- Each unit supports up to 8 simultaneous users
- Remote authentication support, including LDAP, RADIUS, and Active Directory

High Density KVM	Distributed KVM
34 servers may require two 32-port KVMs or some combination that leaves unused ports, increasing the cost per port.	With Spider, you can add one port at a time.
If a 32-port KVM unit fails, all access to those connected servers is lost.	If a unit is disabled you lose access to just one server; and Spider is easily replaced.
If access to all 32 servers is through a single Ethernet connection and the switch port fails, access to all 32 servers is lost.	Spiders can be cascaded or connected individually back to a switch. They can even be distributed among multiple switches so no single switch failure cuts you off from all your servers.
Most high-density solutions require purchasing dongles to connect to managed servers, increasing cost.	Spider has the necessary HD-15 video, USB and PS/2 connections built in.
A typical high-density solution allows access to a limited number of managed servers at a time. (A high-end 32-port KVM may only allow as many as 8 servers to be accessed at once.)	Each Spider allows one or more people to connect to a server. 32 Spiders means 32 simultaneous non-blocking connections - multiple users can access the attached server simultaneously!
Current CAT5-based solutions are limited to 50 to 150 feet between the switch and server. And image quality degrades with distance.	With Spider, there are virtually no cable limitations and image quality doesn't degrade with distance.
Switch-based solutions may require separate KVM and serial dongles and/or external power supplies.	Spider is a completely integrated IP-based KVM switch that does not require an external power supply.

Conclusion

The need for real-time information and online connectivity is at its greatest. With many organizations's network infrastructure spread across corporate campuses, among cities and around the globe, there is a critical market need for the next-generation KVM over IP solution that is specifically designed to meet the needs of the distributed IT environment. The distribution of interconnected resources requires a secure IP-based KVM management strategy that guarantees simple, non-invasive, flexible, fault-tolerant, scalable and cost-effective access… from anywhere, at any time.

When considering a KVM over IP solution, organizations with a distributed IT environment should look for solutions that increase network reliability, improve productivity and reduce hardware/software and maintenance costs. For maintaining 100% network uptime, guaranteed non-blocked access to mission-critical servers is a fundamental requirement from a KVM over IP solution. It also delivers a lower cost per remote user.

To further lower total cost of ownership, built-in Ethernet switching technology, server-powered design and a browser-based interface that does not require any special client software are key features. Finally, scalability and flexibility in a "zero-U" form factor for cost-effectively adding additional servers and/or remote users is an important element for consideration in a distributed IT environment.