Small to Medium Business Solutions

Millions Connected. Millions Saved. It's Your Turn.

XSTACK
Network Storage

D-LINK@WORK
End-to-End Networking Solutions for Your Growing Business

iSCSI Storage for Growing Businesses
Managed SAN Solutions

Solution Brief
September, 2006

Abstract

The D-Link xStack Storage products provide a very cost-effective way for small and medium-sized businesses (SMBs) to meet their growing storage needs. The combination of iSCSI and SATA technology creates an easy to deploy, affordable solution that brings Storage Area Networking within reach of SMBs.

D-Link®
Building Networks for People

Storage Needs for SMBs

Businesses today are faced with an ever-increasing volume of data generated from existing applications like e-mail and ever-larger documents and presentations, along with new sources such as video servers and data warehousing. What was a typical data load only a few years ago has now exploded into levels four times or more in size. Add to these looming future drivers such as government regulatory requirements, corporate disaster recovery planning, etc. and you can quickly see the burgeoning needs of storage will not only grow, but will grow exponentially. Information has indeed become the lifeblood of any company and the crucial challenge for any IT Department is to determine the most effective method to safely store, access and manage this resource.

Many years ago, data was stored on a single server accessed by all employees, but the development of inexpensive machines has meant that nearly every department that asked for its own devoted server soon received it. The consequence of such an environment has been a disconnect between allocated resources and the actual storage needs of any given department. How is an IT manager to control such an environment? Enter the Storage Area Network (SAN) that provides a central location where backup and management of data is made much easier and storage expansion becomes manageable. Most commonly, a high-speed network separate from the Local Area Network (LAN) is created on the "back-end" by interconnecting servers and SAN devices. A better understanding of this is shown in Figure 1. The server storage traffic is isolated within its own separate network where it won't hinder the performance of the "front-end" local area network linking end-user workstations and other devices.

Also integrated into this new paradigm of thinking is the concept of "storage virtualization." Virtualization of storage allows a new approach wherein the data is part of a large central resource pool with no regard for the actual physical location. A server will simply see storage space allocated to it from this pool while the SAN device itself uses hardware and software to map the block-based data to a physical location like a set of hard drives. The file server treats the SAN like it was a local hard drive mounted in the server itself. One of the prime advantages of this storage virtualization is the ease of scalability. Add hard drives to the SAN when needed, and the virtual storage space will grow by the same amount. With this schema, you can quickly see that the network administrator is able to garner greater control of the storage space available to the company. If a department requires extra storage space, then it can quickly be allocated through the SAN's administrative GUI.

A further advantage of SANs and storage virtualization is the increased reliability provided by the ability to easily support modern RAID techniques within the storage array. With direct attached storage (storage tightly coupled with servers), replacing or adding a hard drive typically requires bringing down the server and all of the applications running on it. With a storage array supporting RAID, if a disk fails, the administrator can replace it without the users seeing any service interruption. The same is also true when adding storage resources to an array. The new drives are added to the array, and when they are ready, they are simply added to the pool of virtual storage for later assignment as needed.

The widespread deployment of SANs has been made much simpler with the addition of the familiar Internet Protocol (IP) used to

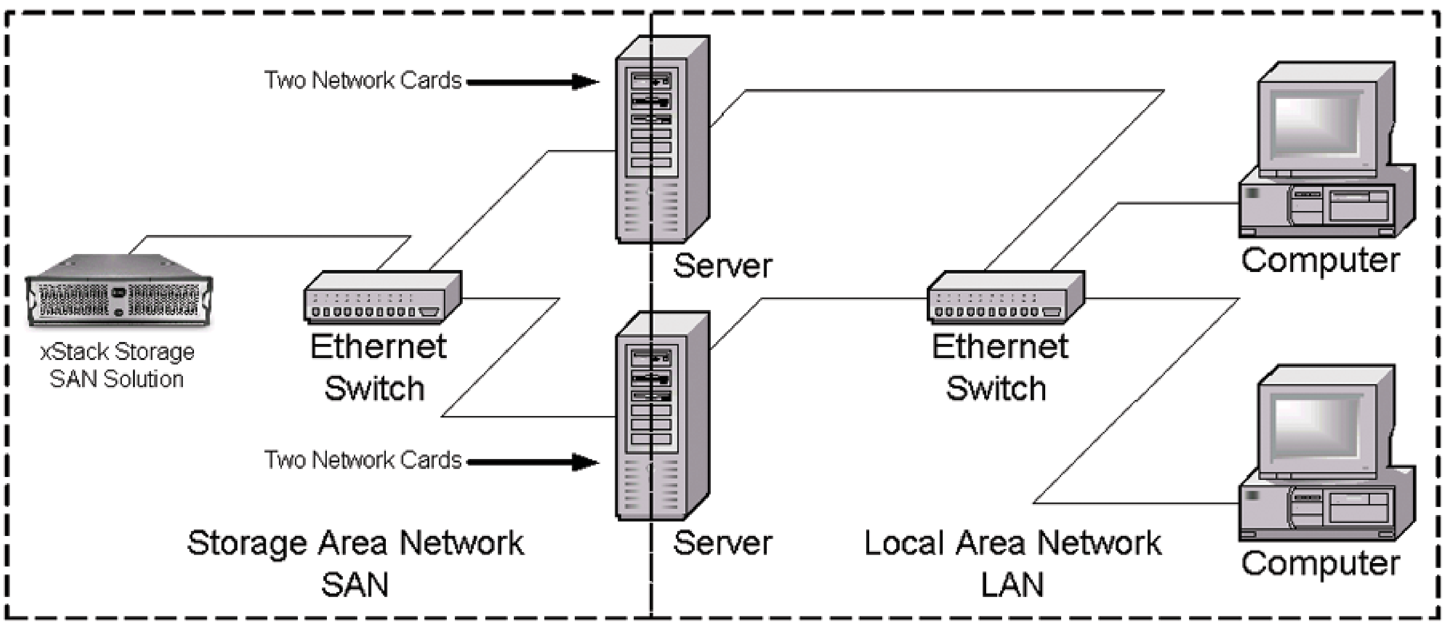

Figure 1: Traditional LAN / SAN Deployment

power the World Wide Web. Converging the world of storage with that of IP networks allows IT managers to create viable SANs using high-speed network equipment that they are already familiar with, such as Ethernet Switches and Routers. There are presently two main standards of SAN technology which encapsulate the storage commands into IP datagrams: Fibre Channel over IP, and iSCSI (IP Small Computer System Interface). Due to the low speeds initially offered by copper-based Ethernet solutions, it was originally passed over in favor of Fibre Channel, which was created as a special dedicated fiber-based system providing speeds of 1 to 2 Gbps (and now 4 Gbps) throughput. As higher speed Ethernet technology became cost effective and widespread, iSCSI has arisen as a viable player for Storage Area Networks. Today, with the promise of 10Gbps Ethernet becoming the high-speed foundation of choice for corporate networks, an IT manager can leverage the department's existing Ethernet backbone to create a fast and efficient iSCSI SAN infrastructure.

An Overview of Storage Area Networking Technology

Even though storage has become more affordable as hard drive prices continue to fall and densities continue to rise, the need to manage this storage and the costs associated with storage management have not decreased. In fact, the combination of new government regulations for storage management, concerns for business disaster planning, and the possibility of a malware attack (virus, hacker, etc.) destroying archives have made automated backups critical to most businesses.

In addition, the constant need for more storage space has become a major element of most storage deployment decisions. System administrators need the ability to easily provide more storage when needed, and manage the storage they have, without bringing down business-critical servers. This is especially critical for many SMBs, who may not have a full-time IT department, or even a dedicated system administrator. When the system administrator is doing the job part-time, the need for a storage solution that is easy to understand and manage becomes even greater.

There are a number of technology alternatives available to meet the storage needs of SMBs. Each of these alternatives have advantages and disadvantages. The following discussion will provide a brief overview of these choices and technologies.

Challenges of Direct Attached Storage

Direct attached storage is generally how most SMBs start out, since servers have drives built in. As the need for increased storage arises, additional hard drives will often be installed directly into the servers. There are a number of problems with this approach, the largest of which is that the server generally has to be taken off-line while the new drive(s) are being installed. In addition, there is a limit to the number of drives that can be supported by a given server. While Direct Attached Storage arrays and servers with RAID support are available, they are more expensive than standard servers, and still have limitations on overall storage size, ability to share the storage across multiple servers, and are time-consuming to manage and backup.

With direct attached storage, managing backups is quite difficult and time consuming. Storage devices are distributed throughout the company, often built into servers and workstations/PCs with different operating systems and usage requirements, making it nearly impossible to create a reliable, automated backup solution.

Another major disadvantage of Direct Attached Storage is the difficulty in utilizing the storage efficiently across multiple servers and users. Drives added to one server are generally not easily available to other servers, so as a company's storage needs grow, management gets increasingly complicated. If a user whose account is on Server1 needs additional storage space, they may not be able to be assigned unused space on Server2 without moving their account. This problem is solved with Storage Virtualization provided by a SAN.

Fibre Channel SANs

The need to address the disadvantages of Direct Attached Storage gave rise to Fibre Channel SANs. Fibre Channel uses the SCSI command set over a dedicated, fiber-based network, to provide very high-performance storage with extensive flexibility for both expansion and distributed storage (the fiber links can span great distances). For many large enterprises, Fibre Channel has been the SAN of choice for business-critical applications that require very high performance (for instance, transaction-based applications). However, this performance, flexibility, and reliability come at a steep price, both in terms of hardware cost and support training costs, and with the arrival of storage arrays based on iSCSI, Fibre Channel is no longer the obvious choice.

Since Fibre Channel uses a dedicated network, it addresses the primary weakness of Direct Attached Storage by separating the servers/users and the storage. This allows much better utilization of the available storage since the pool of available storage can be shared by multiple servers (Storage Virtualization). The pooled storage approach also greatly improves the ease and efficiency of backup operations, as well as expansion and failure replacement. Fibre Channel gets its high performance through a combination of high-speed fiber links (initially 1 and 2 Gbps, now moving to 4 Gbps), and its use of the SCSI command set over a dedicated network. (The SCSI command set is block-based, to directly control the storage devices.)

The primary disadvantage of Fibre Channel is its high cost, both in equipment (Fibre Channel switches, Host Bus Adapters (HBAs), Hard Drives, and the fiber network itself), as well as in the training expenses for the personnel that will administer the SAN.

iSCSI Technology

iSCSI uses the same SCSI command protocol as Fibre Channel, but transports it over standard TCP/IP network (typically using Gigabit or 10 Gigabit Ethernet). Like Fibre Channel, it supports a SAN arrangement, and thus provides the same advantages over Direct Attached Storage. The primary differences between a SAN based on iSCSI and one based on Fibre Channel are cost and ease of support.

The cost savings for an iSCSI based SAN can be dramatic. Not only is the network significantly less expensive due to being based on standard Ethernet switches, but the storage devices and SAN components (hard drives, array controllers, and Host Adapters) are also much less expensive than their Fibre Channel counterparts. In addition, since standard Ethernet devices are used for the network, the amount of training required for support personnel is reduced considerably.

The Host Adapters can be either standard Gigabit or 10 GbE Ethernet devices, or specialized TCP Offload Engines (TOEs). The TOEs implement the TCP/IP protocol to improve the I/O performance of the server by offloading many performance-critical tasks from the CPU.

In the past, FC was the clear leader in performance and reliability, but recently this has begun to change. The performance of a storage array has two components - the actual speed of the hard drives, and the combination of network speed (1 or 10 Gigabit Ethernet vs. 1, 2, or 4 Gbps Fibre Channel), and the overhead added by the TCP/IP protocol. By using Link Aggregation to combine multiple Gigabit Ethernet links into one logical link, an iSCSI SAN array can meet and even exceed the performance of a 2 to 4 Gbps Fibre Channel array, and with 10 GbE interfaces, an iSCSI SAN can support higher interface speeds than Fibre Channel. The difference in drive speed becomes less important when the drives are implemented in a RAID array, since the controller will be accessing multiple drives in parallel, and lower-speed (RPM) SATA drives run cooler, which is a significant advantage in a large array.

The reliability of the hard drives is another area that has seen a dramatic change recently. Fibre Channel drives have always been designed for very high reliability (>1 Million hours MTBF {Mean Time Between Failure}) as well as very high performance. Initially, FC drives had a significant advantage over SATA drives in reliability, but at a steep price premium. Now however, SATA drives are available with similar MTBF numbers (1 Million hours) and the same 5 year warranty as FC drives, but they remain considerably less expensive. With the use of RAID (Redundant Array of Independent Disks) technology in the iSCSI arrays, the overall system reliability for iSCSI can now match that of Fibre Channel.

Hard Drive Technology

There are three main technologies for hard drives at this point:
Fibre Channel (FC): Generally operate at 10,000 or 15,000 rpm, but the drives density has not kept pace with SATA. While they are very fast, they tend to be very expensive, and run hotter than SATA drives.
Serial Attached SCSI (SAS): Performance and reliability similar to FC, but somewhat less expensive. Not yet widely accepted due to SATA's price advantage as an alternative to FC.
Serial ATA (SATA): The higher-speed, serial version of the venerable IDE drive. Newer models are now available with reliability on par with FC, but with higher density (up to 750 GBytes/drive), and lower cost.

All three drive technologies use serial interfaces, and all are hot-swappable. Both Fibre Channel and SAS drives directly support the SCSI command set, while SATA supports the AT command set. For use in iSCSI arrays, the array controller converts the SCSI commands into AT commands for the SATA drives.

Both Fibre Channel (FC) and SAS drives are intended for applications that require extremely high performance and reliability. They typically operate at 10,000 or 15,000 rpm, and are specified with MTBF (Mean Time Between Failure) numbers of 1.2 to 1.4 million hours. SATA drives typically operate at 7,200 rpm, and there are now models with MTBF of 1.2 million hours, and the same 5-year warranty of the FC drives. In a storage array, a robust RAID implementation allows SATA-based arrays to provide performance and reliability previously only available with much more expensive Fibre Channel arrays.

SAS drives have not seen widespread adoption, in spite of their price advantage over Fibre Channel. This is likely due to SATA's significant price advantage over SAS (and even greater price advantage over Fibre Channel). Storage arrays built with SATA drives are significantly less expensive than Fibre Channel arrays, and with a robust RAID implementation, the performance and reliability of a SATA based SAN make it a very attractive choice for many SMBs.

Since the SATA drives are often only a fraction of the cost of equivalent-sized FC drives, multiple spares can be installed into a storage array for rapid recovery in case of a drive failure. In most iSCSI implementations, SATA drives are used in a RAID configuration with one or more "hot spares". Hot Spares are drives that are installed in the array, but not assigned to any RAID group so that if a drive failure occurs, the administrator can quickly assign the spare drive to the group with the failure, thus restoring the data protection. Intelligent storage controllers allow drive reassignment to be done remotely, so that the administrator can simply schedule a time for drive replacement.

With the proper RAID configuration, drive failures can occur without any loss of data or interruption in service to the users. RAID will be discussed in the next section.

RAID Technology

RAID technology (Redundant Array of Independent Disks) is implemented in the disk array controller to allow a set of hard drives to viewed by the network as a group with either higher performance, higher reliability, or just to be seen by the network as a single logical unit {a larger "virtual" drive made up of a number of smaller physical drives}. Some of the popular RAID configurations provide a combination of these three features.

The actual data on the drives is protected by a CRC (cyclic redundancy check) that allows the controller to confirm that the data is not corrupted. RAID techniques are not intended to protect against data corruption, but rather, protect against drive failure. Since hard drives include high-speed mechanical components, they will always have a failure rate. Higher quality drives will have a lower probability of failure, but no drive type is immune to failure. With RAID, (except for RAID 0), the data stored on a RAID set can be recovered after a drive failure (and in most cases, the failed drive is simply replaced with a spare).

There are several common RAID implementations:
JBOD: Just a Bunch Of Disks. Not a RAID mode, but commonly supported to present a concatenated virtual drive to the network using a number of smaller physical drives. May also present the individual drives to the network independently (no concatenation). May be used in combination with one of the other RAID modes.

RAID 0: Striping. RAID 0 is often not considered a true RAID mode since it provides no redundancy. Data is split across two or more drives, without any additional parity information. Generally done to provide improved performance, or to create a large virtual disk out of a number of smaller physical disks. Since the data is spread across multiple drives, the performance improvement ratio should be close to the number of drives in the RAID 0 set (2x for 2 drives, etc.). Both read and write performance is improved since there is no parity calculation being done. If any drive in a RAID 0 set fails, data will be lost, so RAID 0 is rarely used alone.

RAID 1: Mirroring. Data is mirrored (exact copy) across two or more drives. Improves read performance since any copy can be read, and improves reliability dramatically, since a single drive failure does not interrupt operation, or put any data at risk. Since all copies must be written, write performance is not improved over that of a single-disk.

RAID 5: Block-level Striping. Parity information is distributed across the members of the RAID set (generally three or more drives). Popular due to the overall low cost of redundancy. Accesses data in "Stripes", which are a set of data blocks spanning the drives in the RAID set. RAID 5 is more space efficient than RAID 1 since all of the drive space except for the parity portion of each stripe is available for storage. As for performance, RAID 5 suffers for writes that are smaller than a stripe since the controller will have to do a read, modify, write operation to correct the parity block. Read performance can be nearly as high as with RAID 0 by allowing reads of different blocks to be serviced from different drives simultaneously. Since many applications require many more reads than writes, RAID 5 is often a good choice.

RAID 1+0: (Often called RAID 10) A Stripe of Mirrors. Sets of drives are mirrored (RAID 1), and these sets are then striped (RAID 0). Very high performance since parity calculation is not needed. Provides the performance of a striped system, along with the protection of a mirrored system. All but one drive in a given mirror set can fail without losing any data. Since Mirroring is involved, the available storage space will be the size of the mirror.

Modern hardware-based controllers that form the heart of most storage arrays allow complex RAID implementations to operate without a performance penalty. This allows iSCSI users to get the cost benefits of SATA drives while still getting overall array reliability that rivals that of much more expensive Fibre Channel systems.

Since RAID techniques rely on either multiple copies of data, or parity blocks written with the data for protection, drive coherency is critical. If power fails when there are writes in progress, the data on the disk array can be corrupted due to either a parity block that does not correspond to the data in the stripe (partial write), or to a mis-match in a set of mirror copies (unless there are more than two sets in the mirror, in which case, the controller may be able to determine the correct data by using a voting method). For this reason, it is important for Storage Arrays to have a battery-backed cache so that incomplete writes are not lost in the event of a power failure.

The D-Link xStack Storage iSCSI solution

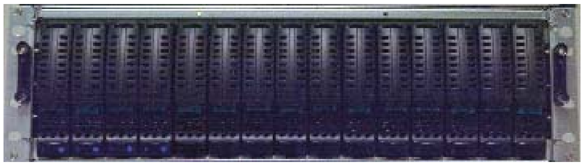

D-Link's xStack Storage solutions leverage the iSCSI standard and SATA hard drives to create a cost-effective group of products designed to meet the needs of a growing company. The DSN-3200 and DSN-3400 products provide over 11 TBytes of storage (using 750 GByte drives), in a compact 3U device. The DSN-3200 provides eight Gigabit Ethernet ports that support IEEE 802.3ad Link Aggregation so that they can provide multi-gigabit links to the network. The DSN-3400 provides a 10Gb Ethernet interface for connection to a high-speed LAN backbone. Both devices feature 15 hot-swappable SATA drives that support RAID 0, 1, 5, and 1+0.

Small business entry level

The D-Link DSN-3200 targets the needs of small companies that have outgrown Direct Attached Storage (DAS) and are now seeking a cost-effective turnkey solution for storage expansion. Providing what is essentially an out-of-the-box SAN, they offer easy configuration and management. Initially, they will support the Microsoft Windows Server family, but will soon be enhanced to address other operating systems such as Linux, Mac OSX and Solaris.

The xStack storage products fit seamlessly into an existing Ethernet LAN, providing familiar management interfaces and are a natural extension of D-Link's existing xStack line of Ethernet switches.

xStack Storage will allow a small business to start with an entry-level system and build incrementally as their storage needs rise. Later as they see their needs increase, they can add drives for additional storage space as needed and additional storage software to meet the requirements of medium-sized businesses. For larger storage needs, multiple xStack Storage devices can be added as needed. As their performance needs increase, the DSN-3400 provides 10G Ethernet to link their storage arrays directly into their high-speed LAN backbone.

Midrange business level

The DSN-3000 families also provide a solid platform for growth by allowing midrange customers to easily expand the unit's functionality. The xStack storage products when used in conjunction with Storage Management software, allow SMBs can implement Enterprise-class storage features at an affordable price. These features include multi-tiered data replication, business continuity and disaster recovery, automated storage snapshots, and centralized storage management and reporting.

DNS-3200/DNS-3400 Description

The D-Link xStack Storage Area Network (SAN) family is designed to provide a powerful data storage solution for small, medium-sized and enterprise customers. At the heart of the DSN-3000 family of products is a 10Gbit iSCSI System-on-a-Chip (SoC) solution that can handle over 65,000 I/Os per second, driving a fully featured storage-virtualization firmware stack. xStack Storage also offers features such as Media Protection, Caching, Redundancy, End-to-End data protection and RAID levels 0, 1, 2, 1+0, 5 and JBOD. A system memory cache of 256MB to 512MB (512MB standard) and a battery-protected cache memory of 256MB to 4GB (512MB standard) allow over 11TB (with 750GB drives) of raw data space to be efficiently accessed. Combining the network and storage subsystems results in a high-performance, scalable, cost-effective solution with state-of the-art reliability.

Competitive products offer "discrete implementations" where a chassis, a main motherboard, a RAID storage controller, network interface cards, operating system software and controlling storage system software are chosen separately and then assembled into a solution. D-Link's xStack can significantly outperform these discrete implementations and does so at a much lower price point by using its highly integrated SoC-based design.

Key Features:

- Supports over 65,000 I/Os per second.
- Embedded disk controller supporting RAID levels 0, 1, 1+0, 5 and JBOD.
- Up to 1,024 virtual volumes with on-line capacity expansion and instant volume access.
- Hot-swappable drive bays support 15 SATA I or SATA II drives with no upper limit on drive capacity.
- Auto-detection of failed drives and auto-rebuild using drive from spare pool.
- Battery-backed cache memory provides data protection even in the event of a power failure.
- Embedded IP-based Management GUI allows remote configuration and monitoring.
- Eight Gigabit Ethernet ports supported on the DSN-3200 provide a high-performance redundant path to the LAN. With Link Aggregation, multiple Gigabit Ethernet ports can be configured as a single logical link to provide multi-gigabit link speeds.
- The DSN-3400 offers a 10GbE full wire-rate iSCSI interface in an XFP connection. (XFP transceiver sold separately).
- Supports CHAP authentication and up to eight VLANs, along with Jumbo Frames.
- Robust (19" rack-mount 3U) chassis is custom designed by

D-Link for strength and durability.
-Drive cage is enclosed with cast aluminum guide plates to prevent flexing and allow easy hard drive insertion and removal.
• Redundant, hot-swappable power supplies provide added system failure protection.
• Enclosure management based on SMI-S v1.1.

xStack Series SAN Arrays

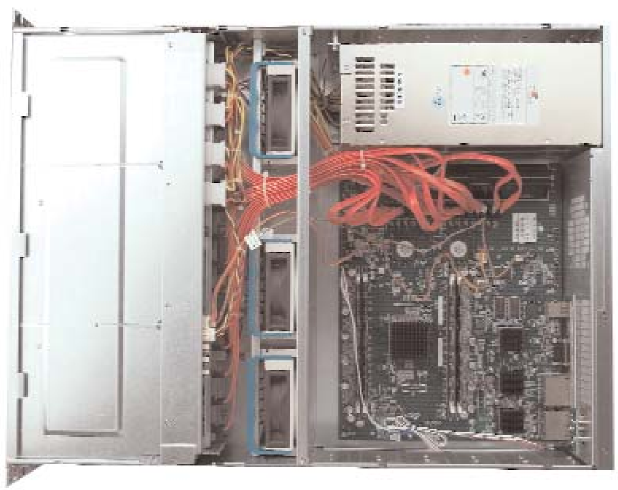

DNS-3200
8 x 1 GbE iSCSI SAN Array
8-Port 10/100/1000BASE-T
Integrated iSCSI System-on-a Chip (SoC)
Handles over 65,000 I/Os per second.
15 hot-swappable SATA I or II drives.

DNS-3400
1 x 10 GbE iSCSI SAN Array
Supports pluggable XFP transceiver
Integrated iSCSI System-on-a Chip (SoC)
Handles over 65,000 I/Os per second.
15 hot-swappable SATA I or II drives.

DEM-421XT
10 Gigabit Ethernet XFP
10GBase-SR
Up to 300 m.

DEM-422XT
10 Gigabit Ethernet XFP
10GBase-LR
Up to 10 km

For more information about products by D-Link please visit the D-Link website at: **http://www.dlink.com/business/**

Or contact D-Link directly:

ASK FOR THE SYSTEMS INTEGRATION TEAM!
Send us an email to **SI@dlink.com**
or call us at **(888) DLINK-SI (888-354-6574).**
Faxing: **(866) 743-4664**
We look forward to building networks with you!