NETGEAR

4500 Great America Parkway • Santa Clara, CA 95054

White Paper

Switch Management: Why Your Small Business Needs It, February 1, 2002

Small-business networking environments are more powerful than ever. From e-mail to office suites and mission-critical applications, networks give employees the tools to do their jobs, while providing small businesses with the performance and flexibility they need to be agile competitors in today's challenging global marketplace.

To remain competitive, small businesses are equipping their employees with desktop personal computers (PCs), notebook PCs, and an Information Technology (IT) infrastructure in an effort to maximize productivity while lowering total operating expenses. In reality, most small businesses with 250 or fewer employees cannot take full advantage of their networked environment unless it can be managed simply, effectively, and affordably. Otherwise, the total costs of ownership (TCO) can spiral out of control, preventing small businesses from realizing the full value of their environment.

Among the hidden costs of a poorly managed small-business network are:

- Lost time due to excessive network downtime.
- Low productivity due to bottlenecks and poor understanding of network traffic patterns.
- Configuration and management tools which are difficult to use.
- Adopting a managed solution that forces customers into buying high-priced products from a single network vendor to ensure interoperability with their existing investment. This limitation prevents customers from mixing and matching best-of-breed products from different vendors, which can offer greater capabilities at a lower cost than a solution they are locked into.

Contrast the points above to a well-managed small-business network that has adopted a managed-switch solution specifically designed for small businesses. In this environment:

- Switch and traffic monitoring help head off problems before they occur, reducing user downtime.
- Management tools offer an intuitive graphical user interface (GUI) that simplifies configuration and monitoring tasks.
- Management functions can be performed remotely using a web browser or directly via a console directly connected to the switch.

- Port trunking is used to aggregate multiple switch ports into one full-duplex, high-bandwidth port for greater throughput.

- The IEEE802.1d Spanning Tree Protocol is used to enhance network resilience to link failures and protect networks against loops.

- Switch firmware can be upgraded through TFTP to take advantage of new features and improvements.

As network-centric applications, file sizes, Internet traffic, network users, and the need for instant access to data all continue to increase, small-businesses will continue to be impacted by the resulting network bandwidth crunch. Because network availability has a dramatic impact on the success of small businesses more and more small businesses will be turning to a managed-switch solution for their networks.

Managed switches designed for small businesses offer plug-and-play installation and easy maintenance, with basic and advanced features that stand up to the bandwidth demands of today. Managed switches also hold the key to the emerging capabilities that small-business networks are expected to provide, now and in the future.

This white paper shows how small businesses can benefit from a managed switch. It describes the basic and advanced management features found in managed switches, and provides practical suggestions for putting managed switches to work for you. A glossary at the end of this paper defines the technical terms used in this paper (these terms are italicized when mentioned for the first time).

What is a Managed Switch?

The network is a vital component of any small business. Without the network, even the best servers, applications, and databases are effectively useless. The one common element of all aspects of any IT infrastructure is the underlying network.

A managed switch allows the ports on the switch to be configured, monitored, enabled, and disabled. Switch management can also gather information on a variety of network parameters, such as:

- The numbers of *packet*s that pass through the switch and each of its ports.
- What types of packets they are.
- Whether the packets contain errors.
- The number of *collision*s that have occurred.

Basic Management Features of a Managed Switch

Switches come with a variety of features. The following sections describe basic features that a small business should look for in a managed switch. These sections include:

- Gigabit Ethernet Support
- Port Monitoring
- Central/Remote Monitoring

- Detecting Packet Errors
- Fault Isolation
- Capacity Planning
- Port Control- Speed, Mode, and Flow Control

Gigabit Ethernet Support

Gigabit Ethernet has emerged as an easy and inexpensive solution for networks straining under the weight of increased network traffic. Gigabit Ethernet offers high-speed 1 Gbps (1000 Mbps) operation, which is 10 times faster than the 100 Mbps speed of Fast Ethernet. All network services and protocols written for 10 and 100 Mbps Ethernet networks work unchanged over Gigabit Ethernet. These similarities allow Gigabit Ethernet to be easily implemented into, and coexist with, existing Ethernet and Fast Ethernet networks. For this reason, small businesses deciding on a managed-switch solution should choose a managed switch that supports Gigabit Ethernet — even if they currently use Ethernet or Fast Ethernet.

Furthermore, small businesses can combine Gigabit Ethernet with other low- or high-bandwidth technologies to structure the environment best suited to their requirements. The following table provides application examples.

Table 1. Examples of Switch Applications Using Gigabit Ethernet

Building backbone	Connect a central Gigabit Ethernet switch with centrally located servers and with workgroup and department switches located in wiring closets.
Campus backbone	Connect building switches to a central campus switch.
High-Performance Servers	Deliver a high-speed connection to powerful computers.
Power Users	Provide high bandwidth to the desktop for multimedia and other data-intensive applications.

Port Monitoring

One of the core objectives of switch management is its ability to monitor the number of data packets sent or received on every switch port. Determining which ports are being used more than others allows companies to identify traffic patterns, monitor network trends, and determine appropriate bandwidth needs so that the appropriate changes can be made to the network to increase performance and efficiency.

By monitoring switch ports continuously, port monitoring also provides companies with the assurance that ports are working properly and are available to users. It also makes it possible to identify port congestion and port failures, minimizing their potential impact on users and the business as a whole.

Some port monitoring features also monitor the types of packets sent and received on switch ports. These data packet types include:

- Unicast packets — data packets with a single destination address. Unicast packets are sent from a single source to a single destination. This node-to-node transmission method is the predominant form of transmission on *Local Area Networks (LANs)* and across the Internet.
- Multicast packets — data packets sent to a selected group of recipients.
- Broadcast packets — data packets sent to everyone on the network.

Central/Remote Port Monitoring

Managed switches can contain embedded software or come with installable software that allows ports to be monitored centrally and/or remotely.

Central port monitoring, as its name implies, provides a central point of control. With this monitoring method, client software is installed on a central workstation or PC. With this software, the health and performance of managed switch ports can be monitored from a central management console, without having to go to the switch's physical location.

Remote port monitoring allows ports to be monitored using a Web browser or SNMP-based solutions:

- Web-based monitoring allows switch ports to be monitored over the Internet. Switches that support web-based monitoring include software for assigning an Internet Protocol (IP) address to the switch. Once the IP address is assigned, you can access the switch from any PC or workstation that has an Internet connection and a Web browser.
- SNMP-based solutions enable industry-standard SNMP and RMON management methods to be used from any network-management application or in-band via Telnet. SNMP and RMON management methods are described later in this paper.

Detecting Packet Errors

In addition to detecting data packets, managed switches can detect and record errors associated with data packets that a port sends or receives. Some network-management tools alert administrators to problems like excessive packet errors or bandwidth overload on a segment. Other tools can perform simple baselining, such as computing packets per second versus collisions per second on a per-port basis.

While the number and types of packet errors that can be tabulated varies between network-management applications, typical errors include:

- Undersize packets — packets smaller than the minimum size defined for IEEE 802.3 packets.
- Oversize packets — packets that exceed the maximum length defined for IEEE 802.3 packets.
- Jabbers — packets that a port receives from a faulty device that transmits oversize packets continuously.
- Fragments — incomplete data packets that a port sends or received.
- Packets that contain a Cyclic Redundancy Check (CRC) or alignment error, which means the data was corrupted during transmission.

Fault Isolation

A fault is a condition that impedes the flow of traffic across the network. Fault isolation, as its name implies, is the ability to locate the source of a failure and isolate it before it impacts network communications.

Managed switches that support fault isolation can identify potential and actual problems caused by various factors, such as traffic patterns, traffic congestion, and port status.

Capacity Planning

Change is the one constant in small-business networking environments, where new applications and new users are continually being added. Despite the dynamic nature of their networks, most small businesses lack the in-house expertise to determine what resources will be required to provide adequate network stability and performance, at a reasonable price, 45, 30, even 15 days from today. Worse, most small businesses can't budget against a guess.

For this reason, some managed switches support capacity planning. These features allow small businesses to obtain a good reading of their current networking requirements so they can ascertain the viability of future solutions.

The objective of capacity planning is to analyze current workloads and reserve capacity, and provide estimates of workload growth and sizing of new applications. Based on that information, estimates of the resources required to support future workloads can be estimated.

With capacity planning, small businesses can:

- **Measure their current network capacity**
  For example, determine which switch ports handle the most network traffic and which handle the least — and at which times.

- **Understand how changes affect their network**
  For example, how adding users and applications affect the network. In some cases, small businesses may need to add more network capacity or upgrade their hardware (such as moving from Ethernet or Fast Ethernet to Gigabit Ethernet) to handle these changes. In other cases, they might be able to accommodate the changes within the current network.

- **Identify network capacity changes**
  As more users and resources are added to the network, small businesses can identify changes and trends in network capacity and utilization patterns and identify network service-improvement opportunities.

- **Plan for growth based on future resource requirements**
  As you become more familiar with your network capacity and the factors that affect growth, and learn how to recognize changes and trends, you can proactively plan for network growth.

In this way, capacity planning allows small businesses to determine where and when to add new capacity, stay ahead of their networking needs, justify purchases more easily, and deploy existing solutions more efficiently.

Port Control - Speed

Managed switch ports support auto-sensing and manual configuration. Auto-sensing enables switch ports to automatically operate at the fastest speed possible, depending on the speed of the attached device. Some newer switches even provide 1000 Mbps (1 Gbps) auto-sensing ports.

Small businesses — especially those on a tight budget — may decide to postpone their purchase of new technologies until the price comes down. Until that time arrives, these businesses will continue using their legacy devices. Certain legacy devices, however, may not support auto-sensing. Auto-sensing switch ports connected to these devices are unable to ascertain the attached device's maximum speed. Therefore, small businesses with legacy devices should look for a managed switch whose port speed can be manually configured to match the speed of the older devices.

Manually configuring a switch's port speed is also helpful when you want to restrict the traffic rate for a specific link to better balance the flow of traffic on the network.

Port Control - Half-/Full-Duplex Mode

Switch ports can use one of two duplex modes:

- Half-duplex — allows packets to be sent or received, but not at the same time.
- Full-duplex — allows packets to be sent and received at the same time. Simultaneous transmission and reception doubles the bandwidth for each port. A 10 or 100 Mbps connection, for example, becomes a 20 or 200 Mbps connection when operating full-duplex.

Most switching ports today support auto-negotiation. Auto-negotiation automatically determines the appropriate duplex mode a port should use, based on the capabilities of the device at the other end of the link.

However, when working with older equipment, the two devices cannot successfully negotiate the fastest connection possible and typically fall back to modes that are less desirable. For this reason, network management allows administrators to manually configure duplex modes for ports.

Port Control - Flow Control

While a high-performance switch forwards data packets at full wire speed to and from its ports simultaneously, there may be times, however, when a switch port may not be able to accept packets at the rate it is receiving them. For example:

- The switch port may be receiving packets from multiple ports at the same time, or
- The switch port may be receiving packets from a port operating at a faster speed. For instance, the sending port might be operating at 100 Mbps, while the receiving port operates at 10 Mbps; or the sending port might operate at 1000 Mbps, while the receiving port operates at 100 or 10 Mbps.

If data packets arrive for a port that is saturated with other packets, the packets may overflow the port's buffer, resulting in dropped packets and lost data.

Flow control is a congestion-control mechanism that prevents data loss at congested ports. Flow control prevents packet loss by controlling the flow of data from the transmitting device to ensure that the receiving device can handle all of the incoming data.

A managed switch typically provides two types of flow control methods for minimizing packet loss: back-pressure flow control and IEEE 802.3x flow control.

- Back-pressure flow control is a collision-based mechanism used with ports operating in half-duplex mode. If the receiving port becomes congested, it fakes a collision with the transmitting port by sending a collision packet to that port. The collision triggers the back-off algorithm defined in the Ethernet specification, causing the sending port to temporarily stop transmitting. This gives the receiving port time to handle the data it has received before processing any more.
- IEEE 802.3 flow control is used with ports operating in full-duplex mode. If the receiving device becomes congested, it sends a pause frame to the transmitting device. The pause frame instructs the transmitting device to stop sending packets for a specific period of time. The transmitting device waits the requested time before sending more data.

Advanced Features of a Managed Switch

The advanced features of a managed switch are designed to optimize total network performance, design, and security. The following sections describe these advanced features:

- Port Trunking
- VLAN support
- Traffic Prioritization and Class of Service (CoS)
- Spanning Tree Support
- Simple Network Management Protocol (SNMP) Management
- Remote Monitoring (RMON)
- Multicasting and IGMP

Port Trunking

Port trunking (or “link aggregation") provides a cost-effective way for small businesses to meet their network capacity and availability needs. It also provides a cost-effective way to increase the bandwidth between switches, or between servers and switches, as your network requirements grow.

With port trunking, multiple switch ports are combined (or aggregated) to form a single high-speed connection. Figure 1 shows an example of port trunking.

With port trunking, bandwidth increases by the number of links combined. For example, using three 100 Mbps connections in a port-trunking configuration results in a 300 Mbps capability. Port trunking also provides redundancy. If one of the combined links fails, the remaining aggregated link(s) continue to function, sharing the traffic for the down link.

In addition, all higher-level network functions — including spanning tree protocol (STP) and virtual LANs (VLANs) — do not distinguish a trunk from any other network port. To these higher-level functions, a trunk is viewed as one big logical link into the network. In this way, these functions treat a trunk as a single entity, rather than as separate links. If you define a

VLAN for a trunk comprised of four aggregated ports, for example, the VLAN treats the trunk as a single link, rather than as four individual connections.

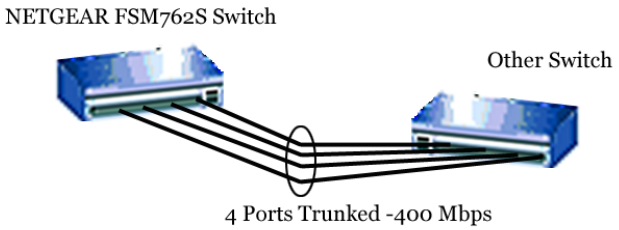

Figure 1. Example of Port Trunking

As an alternative to port trunking, some organizations are turning to Gigabit Ethernet as a solution for handling bandwidth problems.

VLAN Support

A virtual LAN (or VLAN) is a collection of PCs, servers, and other network resources that are grouped together in the same broadcast domain. VLANs allow administrators to segment network traffic into groups, without being restricted by physical connections — a drawback of traditional network design. Because VLANs are software-based, they allow the network structure to quickly and easily adapt to the addition, relocation, or reorganization of nodes, without requiring administrators to touch the hardware or visit the wiring closet.

The freedom afforded by VLANs allows companies to segment their network in ways that make sense for the organization. These include:

- By departmental groups — a company might have one VLAN for the Human Resources department, another for Marketing, and another for Engineering.
- By hierarchical groups — a company might have one VLAN for executives, another for managers, another for general employees, and another for consultants and temporary staff.
- For security purposes — a company might use VLANs to separate departments or systems with sensitive data from the rest of the network to reduce the chance that someone will gain access to information that he or she is not authorized to see.

While you can have more than one VLAN on a switch, computers on different VLANs cannot communicate directly without going through a router (that would defeat the purpose of having a VLAN, which is to isolate a part of the network). Communication between VLANs requires a router. This requirement can enhance network security by encouraging companies to implement standard router-based security measures to restrict access as needed.

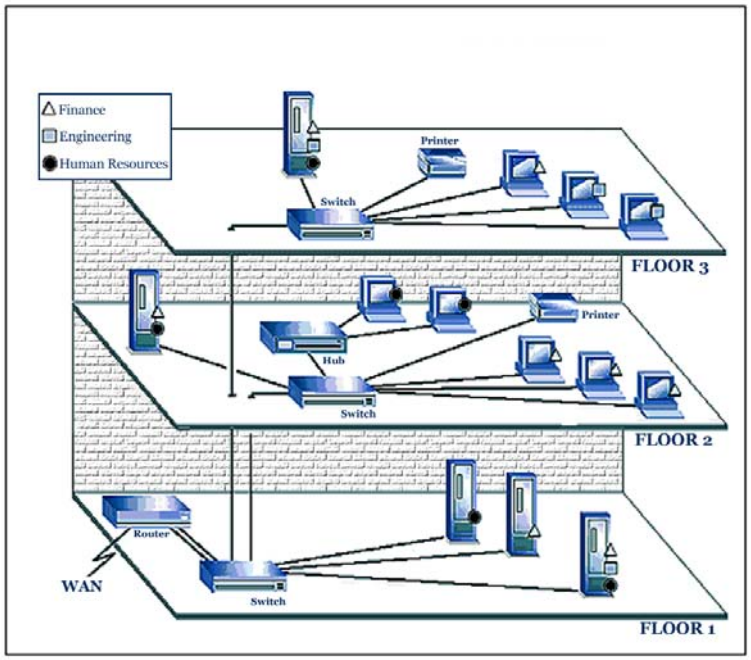

Figure 2. Example of a VLAN

Traffic Prioritization and Class of Service (CoS)

Not all traffic is equal. A financial company, for example, might place greater importance on financial transactions and real-time production data than on other less-critical traffic. Realizing this concern, the IEEE recently ratified its 802.1p specification that defines the standards for traffic prioritization.

Traffic prioritization gives preferential treatment to critical data, enabling the network to speed the delivery of prioritized data to its destination. In fact, the IEEE 802.1p specification defines eight levels of priority to be defined, according to business needs.

Traffic prioritization is especially useful for:

- Converged network applications — organizations that use the same network infrastructure for voice and data may require a higher priority for video and audio traffic to ensure jitter-free video and audio, while relegating Web-browsing video to a lower priority.
- Data-intensive applications — where users need priority connections to server farms and other devices to transfer bandwidth-intensive files.
- Financial applications — where the accounting staff needs instant access to large files and spreadsheets at critical times of the month.

While the IEEE specification has been a long-awaited advancement for switch vendors and users alike, there is a catch: all the applications and devices along a data flow's path must also support CoS to meet the goal of providing end-to-end CoS. This means that, besides switches, devices that must work together can include hosts (clients and servers), routers, and firewalls. Any CoS policy will only be as good as the weakest link in that very lengthy chain. For this reason, complete end-to-end CoS implementations are difficult to achieve and often come at a steep price.

Spanning Tree Support

For an Ethernet network to work properly, only one active path can exist between two stations. A network topology that contains redundant paths creates loops. When loops exist, broadcast packets sent to all switches on the network are flooded back to the sending switch. As the packets are broadcast, received, and rebroadcast by each switch on the network, the number of packets traversing the LAN grows exponentially, creating a broadcast storm that saturates the network and degrades performance.

The spanning tree protocol (STP) enables switches to break loops. Using STP, switch ports identify a single primary path for communicating with a node. Any duplicate paths are placed in standby (or "blocking") mode. Packets are then sent and received through the primary path, creating a loop-free environment. If a change in the network topology renders the primary path unavailable — for example, if a network device fails, is removed, or is added — one of the standby paths is activated.

Simple Network Management Protocol Management

The Simple Network Management Protocol (SNMP) is a set of communication protocols that facilitate the exchange of management information between network devices (such as managed switches). SNMP allows compatible products to be polled from a central management station to gather information on their performance.
SNMP consists of two components: agents and managers.

- An agent is software or firmware included with the SNMP-compliant device. The agent monitors network operations, but does nothing else unless it is polled for information or detects an error.
- A manager is a sophisticated software application that collects and processes information from many agents. It polls certain agents for information at regular intervals or only when requested.

Remote Monitoring

SNMP gathers network data to determine whether SNMP-compliant devices are up or down. Anything regarding traffic management and network performance variations, however, requires additional tools. That's where RMON comes in.

Short for remote monitoring, RMON provides a standard method for monitoring the basic operations of the network. With RMON, administrators can monitor, capture, and analyze real-time information across the entire network using a management station at a central location.

A typical RMON configuration consists of a central network management station, called a probe, and a remote monitoring device, called an RMON agent. As packets travel across the network, an RMON agent installed in a compliant device (such as a managed switch) continuously collects and analyzes network data and stores the statistics locally. This information can be viewed using a probe.

The probe can be a workstation or PC running a network management application, such as SNMPc. From the probe, an administrator can issue SNMP commands requesting information from the RMON agent.

Examples of statistics that an RMON probe can collect include: Bandwidth utilization statistics, number of bytes sent, packets sent and dropped, broadcast and multicast packets, CRC errors, runts, fragments, jabbers, and collisions.

Multicasting and IGMP

The Internet Group Multicast Protocol (IGMP) is a standards-based technology that paves the way for the broadcasting of voice, video, and data over the Internet and corporate *intranet*s, without depleting network resources.

IP multicasting transmits IP packets from one source to many destinations on a network. The network forwards packets to only the hosts that need to receive them. If the network hardware supports multicast, a packet destined for multiple recipients can be sent as a single packet instead of as many copies of the same packet, with each one going to one recipient.

The ability to support thousands of users by using one packet without substantially affecting bandwidth requirements or causing network congestion makes IP multicasting ideal for:

- Distributing internal corporate data to large numbers of users.
- Multicasting live radio and television programs over the Internet, without consuming vast amounts of bandwidth.
- Interactive video conferencing though the Internet, intranet, or extranet, without having to use expensive equipment.
- Delivering internal education simultaneously to hundreds of workers at a single site or at multiple sites.

Conclusion

Networking has emerged as an essential component of small businesses. For many of them, success or failure revolves around the power of the network and their ability to keep it operating at peak efficiency. For this reason, small businesses must manage their networks the same way they manage their other strategic assets.

Rapidly changing traffic patterns, unforeseen network problems, and unexpected killer applications can swamp a small-business network overnight, resulting in network downtime and lost revenue. Since most small businesses have neither the in-house expertise nor personnel to fix network problems, the best strategy is to maximize uptime by protecting networks against potential failures with a managed switch specifically designed for small businesses.

Managed switches aimed at small-business users also offer some or all of the basic and advanced management features described in this paper. Unlike enterprise switches that provide complicated management features designed for corporate networks, the management features offered by small-business switches are designed to be as easy to use as possible. After all, in a small-business environment, a management feature's ease of use can ultimately determine whether that feature is used. The fact that you may be using the switch in the heat of a network outage makes ease of use as — or even more — paramount as the features it supports.

Managed switches designed for small businesses are cost-conscious solutions that offer easy installation, affordable high-speed connectivity, and simplified maintenance, along with the

scalability to support a business' current and future needs. This approach allows small businesses that are building their networks with future growth in mind to start small while thinking big.

Glossary

This glossary defines the technical terms used in this white paper.

Term	Definition
Collision	The situation where two or more stations try to send a data packet on the same network at the same time. The result of a collision is generally a garbled message.
Internet	A global network connecting millions of computers.
Intranet	A network based on TCP/IP protocols that belongs to an organization, usually a corporation, and is accessible only by the organization's members, employees, or others with authorization. An intranet's Web sites look and act like any other Web sites, but the firewall surrounding an intranet fends off unauthorized access.
Local Area Network (LAN)	A computer network that spans a relatively small area. Most LANs are confined to a single building or group of buildings.
Packet	A piece of a message transmitted over a packet-switching network. Besides data, a packet contains the address of its destination.
Segment	A network that is bounded by switches. Dividing an Ethernet into multiple segments is a common way to increase bandwidth on the LAN.

Appendix

NETGEAR's FSM726S Managed Stackable Switch was designed for the small business. Combining a powerful hardware solution with easy-to-use software, this switch is ideal for the first time management user. With 24 10/100 Mbps ports and two, gigabit ports and two stacking ports, the FSM726S can stack to support 144 10/100 Mbps ports and 12 gigabit ports. More importantly, the FSM726S supports the basic and advanced management features described in this paper, enabling businesses start with basic switch management and safely grow with their investment to the advanced features:

- **Intuitive management features**
  The FSM726S can be monitored and managed from anywhere on the network using a simple Web browser. The browser Graphical User Interface (GUI) enables the basic features of:
  - **Port Monitoring**
  - **Detection of Packet Errors**
  - **Fault Isolation**
  - **Capacity Planning**
  - **Port Speed and mode control**

For those users ready for advanced features, the FSM726S also supports:

- **Port trunking**
  Allows multiple ports to be combined into a single high-speed connection to increase throughput and resiliency.
- **VLAN tagging**
  Allows logical groups of users to be defined, without regard to the location of their physical connections to the network. This functionality eases administrative burden of adds, moves, and changes, and confines broadcast traffic within the boundaries of the VLAN.
- **CoS features**
  CoS features allow time-sensitive data traffic to be given a higher priority over less-sensitive data traffic.
- **Spanning Tree**
  Allows the switch to dynamically monitor its topology and prevent multiple paths a loop condition) that could disable a network.
- **SNMP Support**
  Enables management via standard protocols to work with Network Management Software.
- **IGMP snooping**
  Optimizes the flow of IP multicast traffic, such as video and audio broadcasts, enabling advanced multimedia applications to be delivered easily to their destination.

The FSM726S enables small businesses to enjoy the benefits of managed networking with a product designed specifically for their needs.