

Purpose-Built S60 Optimized for High Performance Storage

Storage networks present major challenges for data center switches that must be resolved to ensure fast and error-free transport of storage traffic. Force10 S60 switches are purpose-built to handle storage traffic with line-rate performance in Gigabit Ethernet and 10-Gigabit Ethernet networks.

Storage Traffic Bottlenecks

Storage networks have several characteristics that cause bottlenecks in traffic flow, which reduce performance and cause packet retransmissions. The Force10 S60 switch addresses these bottlenecks.

- 1. Bandwidth Mismatch:** Storage networks are migrating to 10-Gigabit Ethernet, while servers still provide 1G connectivity, which creates a mismatch in bandwidth when the storage traffic enters Gigabit Ethernet ports on a switch. The solution is to use a very large buffer. Unfortunately, most network switches provide only 8 to 16 megabytes of buffer, which is not enough to handle traffic spikes.
- 2. Bursty Application Traffic:** Asynchronous, bursty applications cause network congestion and packet loss. The S60, with the industry's largest packet buffer (1.25 Gigabytes), optimizes traffic buffering for specific application needs, delivers low application latency, maintains predictable network performance and helps deliver required QoS.
- 3. Traffic Spikes:** Traffic spikes occur in a number of scenarios. For example, when transferring large amounts of data to a remote location, retransmission operations are likely to compromise replication performance. Storage traffic is typically carried over TCP-IP transport protocols, which are very sensitive to packet loss. The protocol will "back off," effectively lowering the transmission rate as the window size is adjusted or session restart takes place, which will cause packet retransmission. Large buffers in the S60 minimize or remove any packet loss in the network, which optimizes client, network and storage performance.
- 4. Consolidating LAN & SAN:** Another challenge in storage applications is the emerging trend toward consolidating server and storage traffic on the same network using protocols such as iSCSI and Fibre Channel over Ethernet (FCoE). These consolidated networks must be immune to packet loss. Again, Force10's S60 switch handles consolidated traffic loads with the industry's largest packet buffer.

DCB (Data Center Bridging) is also being explored as an alternative solution. However, this solution is not being considered for 1G to 10G bandwidth mismatch.
- 5. High Performance Audio/Video Applications:** Audio/video applications require high performance access for reading, writing, and editing. Upon access for editing or streaming, the user experience needs to be reliable with no interruption. Over time, data grows exponentially, although performance must stay at the highest level. The key to delivering optimized results is to provide scalability. As nodes are added, performance increases in a linear fashion. This is true for both compute nodes and storage nodes. Force10 S60 switches are ideally suited for this application since the large buffers reduce congestion – eliminating at best or minimizing at worst – packet retransmits, ensuring predictable, consistent network performance.
- 6. Memcache bottlenecks:** In the interest of personalized services, web portals transitioned to using web pages that are database-driven, dynamically generated on the fly. This database load led to the development of a dynamic object-caching layer – the memcache, to extend the scalability of the databases. Objects stored in memcache are typically the results of database queries and it is common for certain objects to get 'hot' and get flooded with microbursts of traffic requests and subsequent responses. These microbursts often overflow switch buffers, forcing retransmits. The S60 alleviates this problem with deep packet buffers, improving application throughput by eliminating the need to re-request packets lost due to buffer overflows.

Business Impacts of Storage Bottlenecks

Storage bottlenecks not only impact the network, they impact business operations.

1. **Revenue impact:** Slower web response in Web 2.0, eCommerce, cloud and similar applications impacts business, thus creating revenue challenges.
2. **Increased costs:** Inefficient architectures that increase the complexity with larger numbers of storage and compute nodes create bottlenecks, especially in cluster, grid and cloud computing environments. This could be addressed more effectively with deep buffer switches.
3. **Slow batch processing:** Large quarter-end batch processing operations on applications like Oracle Financials general ledger closings stress network bandwidth, impacting business productivity and reporting conformance to SEC guidelines.
4. **Slow, error-prone data collection:** Scientific applications that process massive amounts of data create buffer overruns which cause dropped packets, impacting timely completion of scientific research.
5. **Longer backup windows:** When bottlenecks slow traffic to the point that it is not possible to complete a backup in the allotted time the business suffers.
6. **Motion video jitter:** In media editing, video surveillance and video provisioning overloaded packet buffers cause jitter impacting productivity, viewer experience and service level degradation.

The Force10 Solution

Force 10's S60 ToR switch is purpose-built for storage applications with 1.2 Gigabytes of tunable buffers epitomizing data throughput, complementing a storage array's caching functions.

The S60 also supports bare metal provisioning, enabling users to configure hundreds of switches within hours. Force10's JumpStart allows administrators to write simple Perl and Python scripts to develop a master switch configuration and automate provisioning, eliminating potential human errors and ensuring consistency while facilitating switch deployments in few minutes. A tier 1 Web 2.0 gaming company used JumpStart to provision several hundred switches simultaneously.

Use Cases

Several Force10 customers attribute S60's success in their storage environments to its deep buffers and predictable, consistent, low latencies.

1. **myYearbook**, a top-tier social networking provider, uses S60 ToR switches for their deep buffers and uses S4810 switches to face 10GbE storage, linking the two via 10GbE ports. This combination enabled them to implement a video chat service for 25 million users.
2. **An Enterprise customer** uses 1-Gigabit servers and NetApp 3210 storage with S60s, providing the connectivity to run Oracle applications, Sharepoint, web, and other customer-facing applications. They find immense value in their investment protection in servers when they moved to higher speed storage and took advantage of 40 Gbps connectivity.
3. **The High Definition services lab** for a major U.S. semiconductor manufacturer and player in the mobile phone market uses the S60 for one of the most challenging multicast environments. The customer's application uses the S60's deep buffers along with a robust multicast application for processing SD and HD video streams. The application accepts Variable Bit Rate HD video in MPEG-2 format and performs statistical remultiplexing, rate shaping and localized digital advertisement insertion. This video is then streamed to the set-top boxes to subscriber homes. S60 performs exceptionally well in handling thousands of multicast groups and the variable bit rate at which this traffic streams. The result is crystal clear, pixel-free HD video.
4. **A large European Nuclear Research Organization** uses S60 for particle collider experiments. Nuclear experiments generate huge amounts of data from sensors connected to a collider. The sensors use 1Gb Ethernet connectivity to S60. The traffic profile is very bursty and packet drops are very expensive. Thanks to the S60's large buffers and the buffer tuning mechanisms in FTOS, the S60 switch operating system ensures no packet loss. The 1G connectivity, large buffers and low jitter enable the European organization to concentrate on their experiments, increasing productivity.

5. A **media company** uses S60s for streaming large files to its team members for video editing. The overall team productivity skyrocketed with S60s bridging servers and end points to the storage subsystems. The S60's deep buffers help the team with fast access to relevant clips, while its low latency and jitter ensure a smooth playback response.

Non-linear editing for films and television postproduction is a modern editing method, which enables direct access to any frame in a digital video clip, without needing to play or scrub/shuttle through adjacent footage to reach it, as was necessary with historical videotape editing systems. It is the most natural approach when all assets are available as files on disks rather than recordings on reels or tapes. The faster, more deterministic network performance that the S60 provides enables video professionals to "cut and paste" films, increasing their productivity many fold.

Video and audio data are first captured to hard disks or other digital storage devices. The data is either recorded directly to the storage device or is imported from another source. Once imported, they can be edited on a computer using a wide range of software.

Good network performance makes it possible to edit both standard-definition broadcast quality and high definition broadcast quality very quickly and do the full processing of the huge full-quality high-resolution data in real-time.

6. A **video surveillance application** uses S60 switches to link feeds from remote video cameras to the IP network, ensuring errorless, jitter-free video.

An Internet protocol camera, or IP camera, records directly to digital storage media such as flash drives, hard disk drives or network attached storage and sends and receives data via a computer network and the Internet.

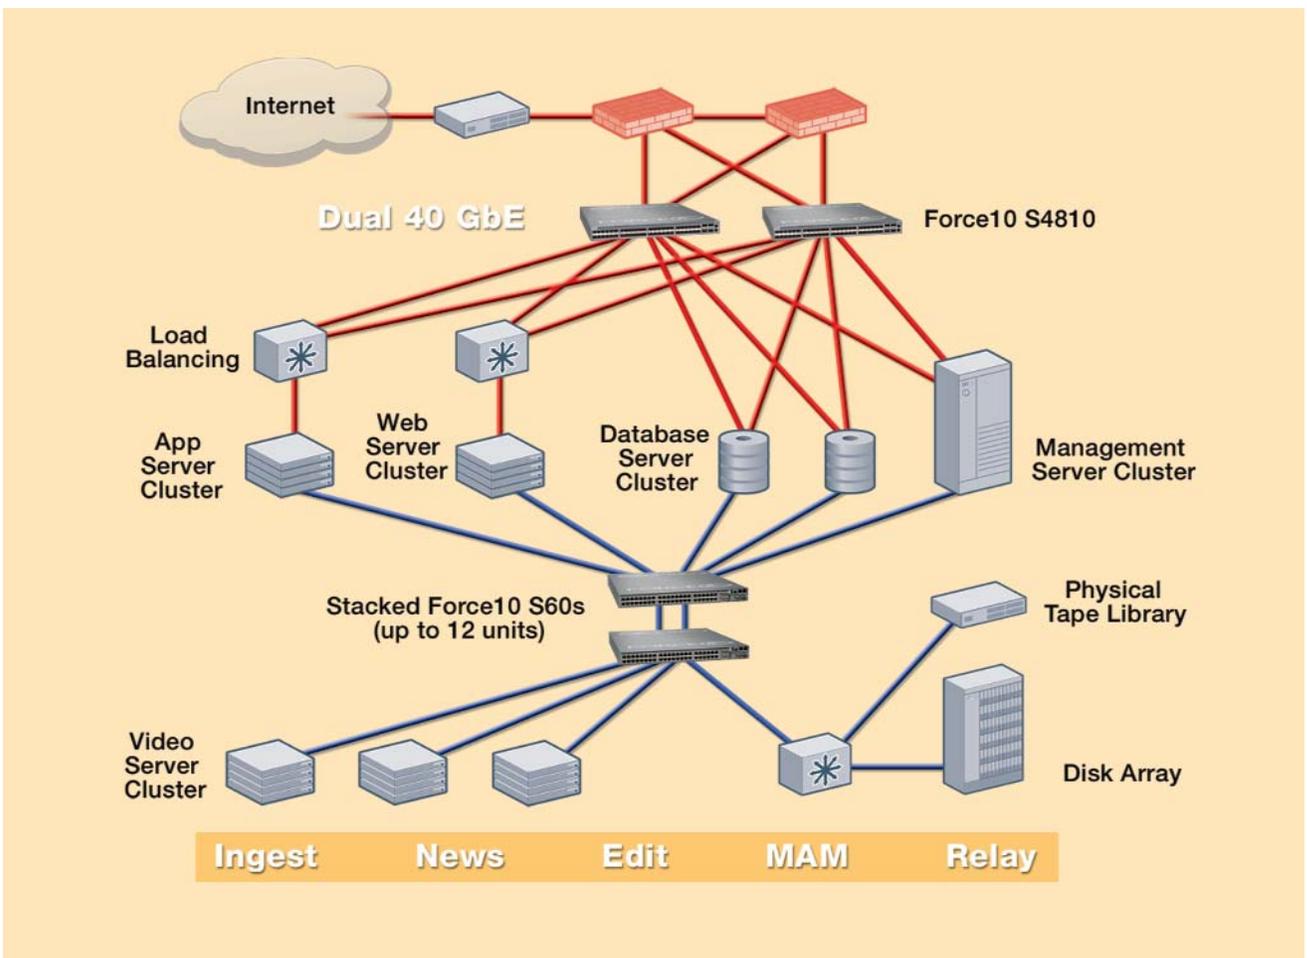


Figure 1. Media Production

Video surveillance is widely in use in banks, casinos, airports, military installations, and convenience stores.

A major Midwestern city operates a networked video surveillance system which combines CCTV (Closed Circuit Television) video feeds of government agencies with those of the private sector, installed in city buses, businesses, public schools, subway stations, housing projects, etc. Even home owners are able to contribute footage. It is estimated to incorporate the video feeds of a total of 15,000 cameras.

The city's Office of Emergency Management detects the caller's location and instantly displays the real-time video feed of the nearest security camera to the operator, not requiring any user intervention. While the system is far too vast to allow complete real-time monitoring, it stores the video data for later usage in order to provide possible evidence in criminal cases.

A major European capital city also has a network of CCTV systems that allows multiple authorities to view and control CCTV cameras in real time. The system allows authorities including the Metropolitan Police Service, Transport for the city and a number of its

boroughs to share CCTV images between them. It uses a network protocol called Television Network Protocol to allow access to many more cameras than each individual system owner could afford to run and maintain.

- 7. **Several HPC clusters** rely on S60 switches to connect Gigabit Ethernet servers with a storage network, and the S60's large packet buffers ensure that all data is sent to and received from storage with no re-transmits.

The S60 has additional features beyond deep packet buffers that facilitate its use and reduce operating expenses. Force10's Open Automation features reduce setup and configuration time to deliver major savings.

- 8. **A Web 2.0 / Portal company** provisioned 300 S60 switches to upgrade their storage infrastructure while protecting its investment in their 1-Gigabit server infrastructure and realized intended performance gains. Force10's bare metal provisioning (JumpStart) feature allowed the company to configure 75 racks in 3.5 hours. The company estimated it could deploy 1.2 MW of data center capacity in a day. Assuming all of the switches are set up to use JumpStart, it is a simple matter of updating the DHCP entry used by the

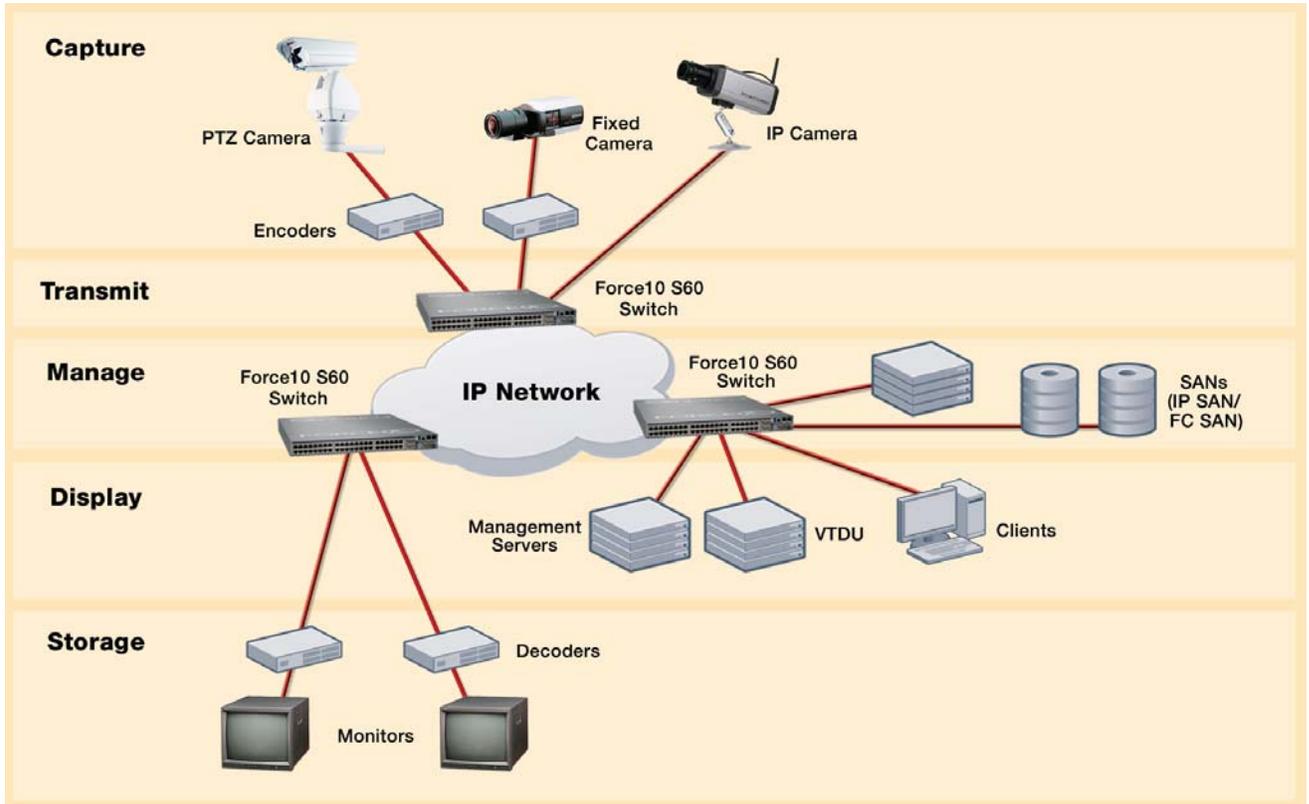


Figure 2. Video Surveillance

switches for the 'current FTOS version.' Then each switch is reset on a convenient timeline and it will automatically upgrade itself during the boot process. This eliminates manual provisioning and the potential for keyboard errors during that process.

Summary

Storage traffic bursts create significant bottlenecks in the network due to - inadequate buffer sizes of most ToR switches. Force10's S60 features 1.28 Gigabytes, the industry's largest network switch buffer that complements

storage caching algorithms and mechanisms. From faster backups and replications to end-of-month closings, Force10 S60 switches significantly reduce the time it takes to complete tasks and drive higher performance eliminating the need to add additional servers and storage systems to address erroneous network bottlenecks, contributing to increased cost and complexity. In addition, the S60's JumpStart provisioning feature makes it possible to provision hundreds of switches in a very short time, with pre-determined, customer specific configurations.

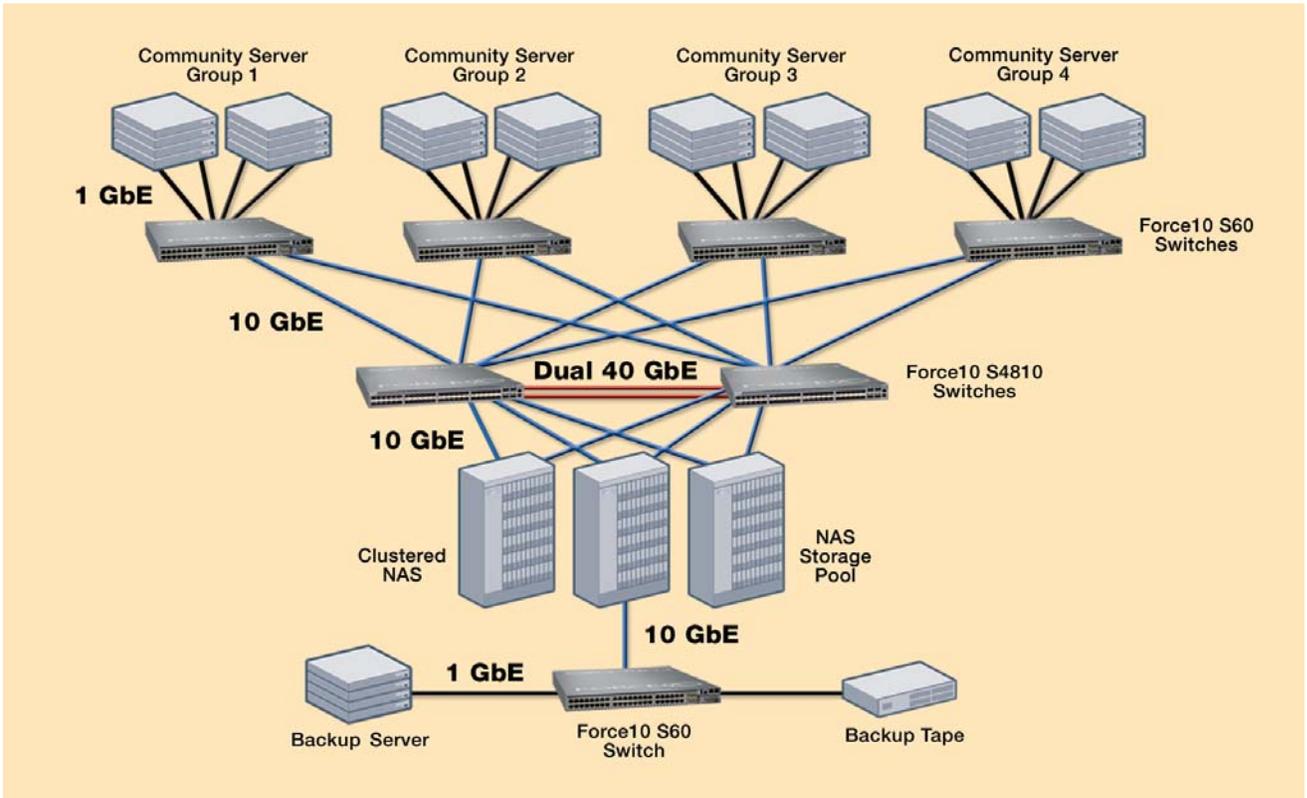


Figure 3. High Performance Computing