

**An Innovative, Ethernet-Based Communications Network Architecture for Integrated
Transportation Management Systems**

Douglas Gettman, Ph.D. (corresponding author)¹
Siemens Energy & Automation – Gardner Transportation Systems
1355 Willow Way #110
Concord, CA 94520
(925) 691-9524
Douglas.Gettman@gts.sea.siemens.com

David O’Keeffe
City and County of San Francisco
Department of Parking and Traffic
25 Van Ness Avenue, Suite 210
San Francisco, CA 94102
(415) 554-2314
david_o’keeffe@ci.sf.ca.us

Kevin Aguigui, P.E.
DKS Associates
1956 Webster Street, Suite 300
Oakland, CA 94612
(510) 763-2061
kga@dksassociates.com

Warren Tighe, P.E. and Diederick Van Dillen
Siemens Energy & Automation – Gardner Transportation Systems
1355 Willow Way #110
Concord, CA 94520
(925) 691-9524
warren.tighe@gts.sea.siemens.com
diederick.vandillen@gts.sea.siemens.com

Bill Cormier and Gordon Jennings
Extreme Networks
3585 Monroe Street
Santa Clara, CA 95051
Gjennings@extremenetworks.com
Bcormier@extremenetworks.com

Abstract

This paper describes an innovative, end-to-end Ethernet-based communications network architecture designed for highly reliable, fault-tolerant communications between a central transportation management center (TMC) and field devices (traffic signal controllers, closed circuit television (CCTV) cameras, changeable message signs, etc.). The concept of operations driving the development of the architecture is the Integrated Transportation Management System currently being developed and implemented for the City and County of San Francisco by DKS Associates and Siemens-GTS.

The architecture consists of three-levels with multiple degrees of redundancy to ensure continuous communications between field devices and the TMC in an environment (i.e. city streets) where accidental cuts to underground infrastructure are routine. At the first level, each field cabinet houses a field-hardened Ethernet switch that connects the field devices to the communications system. The field-hardened Ethernet switches are then connected to two geographically separate distribution switches via single-mode fiber – one as the primary communications link to the TMC and one as a backup. These distribution (“middle”) switches constitute the second level of the communications architecture. The design calls for middle switches operating at wire-speed with the capacity to handle full non-blocking communications for up to 40 connections (ports) as well as all of the network data from a lateral connection to an identical middle switch. Each middle switch is also required to have full non-blocking forwarding of video feeds of up to 320 cameras at 3mbps quality using MPEG2.

Middle switches are placed in the field in climate-controlled hub cabinets and linked back to a central switch located at the TMC. A back up central switch is located at a separate facility. The links between the field hubs and the TMC operate at gigabit speed. The central switch at the TMC is an enterprise-class device with expansion capabilities up to 64 gigabit ports and wire-speed backplane. Fault tolerance and data reliability is provided by standard Quality of Service (IEEE 802.1P) and Spanning Tree (IEEE 802.1D) Ethernet protocols. The architecture was tested and demonstrated at the Extreme Networks Interoperability Laboratory in Santa Clara, CA and was found to be a viable alternative to ATM, SONET, frame relay, and other networking technologies.

Concept of Operation

Many hundreds of color video cameras will be used for monitoring transportation conditions throughout the City of San Francisco. Some are moveable (pan/tilt/zoom) cameras and others are stationary cameras with fixed views. The cameras are connected to the TMC by a communications network with enough bandwidth (except under extreme multiple failure conditions) to bring *all* camera feeds back to the TMC simultaneously, 24 hours a day, 7 days a week. All cameras transmit a broadcast-quality video stream to the TMC. Lower-quality video is available for distribution to partner agencies, the internet, and intranet users. Traffic controllers, changeable message signs (CMS), and other field devices are controlled and monitored through the same communications network as the video. Although a single video stream is roughly 300 times larger than a data stream for a traffic signal controller or CMS, all monitoring, status, and control data (including camera pan-tilt-zoom commands) are not interrupted by video communications. In the case of extreme link failures of the network, some video streams are not available at the TMC and others are available at full-quality depending on the assignment of quality of service.

In addition to operators in the TMC and San Francisco Department of Parking and Traffic (DPT) personnel on the TMC intranet, DPT personnel in the field have the capability to connect a laptop computer to a device in a field cabinet and fully access all functions available at the TMC, including viewing an unlimited number of video camera feeds. A second location (satellite TMC) replicates all functionality of the primary TMC in an emergency situation, such as a power failure at the main TMC location. The status of the communications network is observable from the TMC, and any computer on the DPT intranet with the required software. The communications network is *automatically* fault-tolerant to some degree of device and transmission media failure and can be *manually* re-configured in the event of other failure conditions. Communications media, protocols, devices, and equipment are based on standard to the largest extent possible for future growth. Dark fibers are provided wherever possible for future expansion and to support peer-to-peer communications (if required) for adaptive control and other applications or uses (including non-transportation related uses).

Description of the Ethernet Communications Network

In general, the Ethernet solution is composed of two elements:

- (1) Transmission media/connection layout comprising the *links* of the network, and
- (2) A network of switches, comprising the *nodes* of the network.

Transmission Media

The size of the SF ITMS communications system will be extensive when fully implemented. The number of field devices will possibly exceed 2,000 and the cable routing will be complex. All connections in the communications network are made with single-mode fiber-optic cable (except for connections between the field devices and the edge

switch, which are connected using Category 5 cable). Single-mode fiber optic cable can transmit signals at distances approaching 40km, with appropriate transceiver hardware. Connections to “isolated” traffic signals and devices other than CCTV cameras (i.e. CMS) that have low bandwidth requirements can easily be over wireless connections, (e.g. CDPD). It is possible, also, that in the future, wireless technologies such as CDPD will support data rates that can provide adequate-quality video transmission, but such technology is not currently generally available.

Single-mode fiber optic cable provides the best option for long-term support of the SF ITMS, primarily for provision of “virtually unlimited” bandwidth and the distances that must be implemented for many of the links in the network. Both multi-mode fiber and copper media were considered for short-range links, but neither are recommended for implementation. Initially, it was thought that connections could possibly be implemented between intersection controllers and adjacent controller cabinets using Category 5 (Cat5 or Cat5e or higher-grade types) cable for many of the intersections in San Francisco, primarily in the very dense downtown area, but simple bench testing in the Siemens GTS lab indicated that the IEEE 802-3 distance standard of 100m (320 feet) could not be extended to 400 feet, the assumed minimum possible length to connect the closest intersections in the San Francisco network.

In addition to the technical constraints listed above, direct single-mode fiber output will likely be supported by future 2070 hardware (and likely all other field devices) and as such any marginal cost savings due to the selection of lower-cost multi-mode or copper connections to each individual intersection will be outweighed by the additional issues involved with maintaining two media types, such as additional maintenance, tracking, training, and assessment/error detection issues. For example, copper to *multi-mode* fiber media converters are 30%-40% less expensive than copper to *single-mode* fiber media converters as well as switch ports for multi-mode fiber are 30-40% less expensive than a single-mode port costs. However, using two types of fiber-optics media only adds unnecessary complexity to an already complex system.

Network of Switches

The network of switches performs three primary functions:

- (1) “Aggregation” of lower-bandwidth data streams into larger-bandwidth data streams,
- (2) Routing of individual data packets to the correct destination using the most efficient and currently available route,
- (3) Automatic handling of fiber cuts and switch failure.

Ethernet provides a minimum of 10 Megabits-per-second (10Mbps) communications to each device cabinet, which is more than sufficient for the data rates of traffic signal controllers, CMSs, and most other “data” needs. The driving need for high-bandwidth communications is high-quality video transport. Digital compressed video of DVD quality (MPEG-2 compression) requires bandwidth of at least 6Mbps, which can easily be handled by a 10Mbps connection. Multiple video streams of 6MBps can be aggregated onto a

single 100Mbps connection, assuming a safety factor of not aggregating lower-bandwidth streams up to the maximum capacity of the higher bandwidth connection. We assume 8-10Mbps connections can be aggregated onto a single 100Mbps connection when the switch handling the incoming traffic has a full wire-speed backplane and traffic is routed to a 1Gbps uplink connection back to the TMC. Recall that the utilization of the 10Mbps connection is 6Mbps (plus some background “noise” consisting of traffic signal controller commands, PTZ commands, etc.), so the loading on the 100Mbps connection is approximately 48Mbps (50%). Multiple 100Mbps connections can then be aggregated into 1-gigabit (1000Mbps) bandwidth data streams with a similar safety factor (which, as was shown in network testing, is not needed). Note that all connections are full-duplex connections.

Logical Network Architecture

The Ethernet solution architecture is illustrated in Figure 1 below. The architecture consists of three levels of switching and manual patch panels at cabinets not containing any CCTV cameras. At field locations with more than one connected device, traffic controllers and CCTV cameras (and other field devices such as CMS) are connected to a field-hardened edge switch via Category 5 cable at 10Mbps (full-duplex). Note that 2070 controllers can only communicate at 10Mbps.

The edge switch then aggregates those 10Mbps Ethernet connections to a 100Mbps full-duplex Ethernet connection on single-mode fiber. Two single-mode fiber outputs from the edge switch are connected to a single-mode fiber patch panel and the patch panel is then connected to two middle switches. It is significant to note that the edge switches are not “daisy-chained” – the architecture requires large fiber counts in some conduits (50+ in some street configurations – and the main “trunk” cables have been specified at 144 fiber counts). Compared to the cost of providing new conduit, large count fiber cable is relatively inexpensive.

One middle switch is used for uplink traffic under normal conditions and the second connection to a different middle switch is used as a *backup* connection in the event of failure of the primary middle switch or a cut to the primary connection. Field cabinet locations with only a 2070 traffic signal controller have a media converter to convert Cat5 cable to single mode fiber (instead of an edge switch) that is connected to the patch panel. That same patch panel is then connected to *two* middle switches, for the same reason that the edge switches are connected to two middle switches. In the event of failure of the primary link, a signal technician can manually re-connect the controller to the other patch panel connection.

San Francisco Integrated Transportation Management System – Communications System Architecture

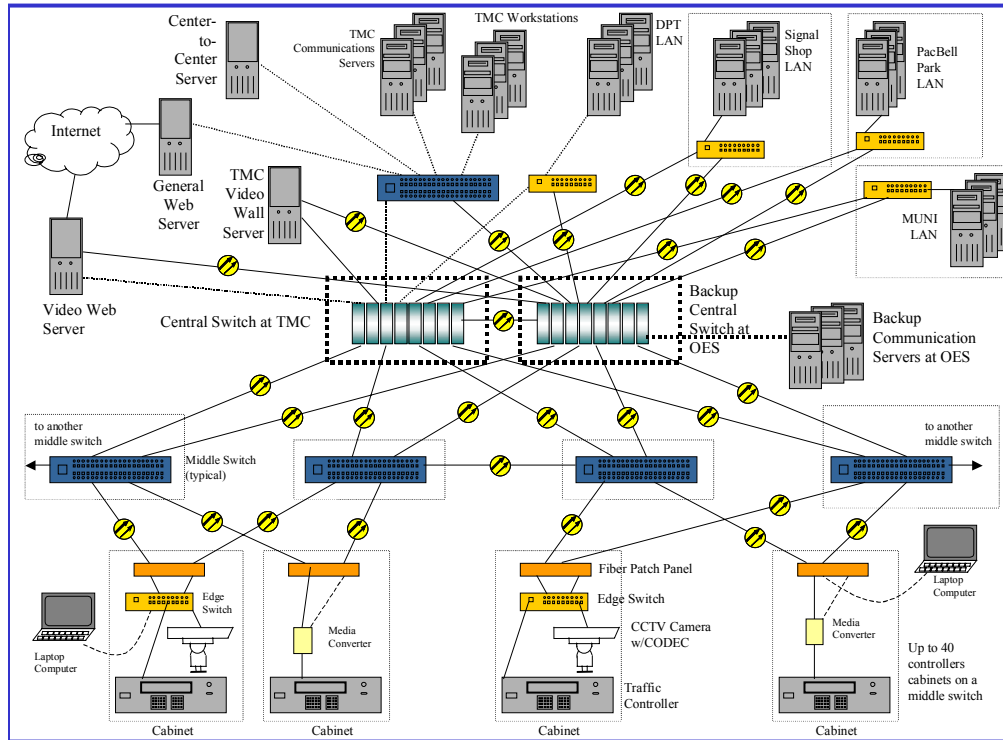


Figure 1: Communications Network Architecture

Each middle switch is connected to *one* other middle switch as a backup route with a 1Gbps connection for traffic in the event that both uplinks from a given middle switch are cut or fail. The architecture design calls for this 1Gbps cross-link to be bandwidth-limited to the maximum loading of the middle switch's primary connections and the edge switches that are connected to this middle switch as their backup connection.

This implies that each middle switch may be required to support:

- (1) The traffic to and from its primary connections,
- (2) The traffic to and from its edge switches and field devices that are using this switch as the "backup", and
- (3) All of the traffic from its *partner* middle switch (i.e. the partner's primary and backup field connections).

As shown in the logical architecture diagram, middle switches that are paired as partners in this way do *not* share the same set of primary and backup connections. A link between middle switches that share backup and primary edge switches is not a useful connection for redundancy requirements. In the field, that link will likely be in the same conduit used to connect the primary and backup cabinets to each middle switch cabinet.

Although it is shown in the logical figure that pairs of middle switches share the *same* sets of primary and backup cabinet connections (i.e. the primary middle switch for one set of

field cabinets is the backup middle switch for the second set of field cabinets, and vice versa), this may not be the case throughout the City. In fact, it is more probable that:

- (1) Field cabinets are grouped geographically and routed to middle switches for *primary* connections as a *group*, and
- (2) Field cabinets are assigned to backup middle switches as an *individual* cabinet-by-cabinet decision based on the fiber distance and number of splices required to create the backup connection.

Although the intent of the logical architecture is consistency of deployment throughout the City, because of the phased implementation of the SF ITMS program and the high cost of conduit installation, primary connections are considered first and redundancy is acquired/provided as it becomes available (i.e. as loops are formed when conduit is installed).

Each middle switch is connected to:

- (a) The central switch at the TMC, and
- (b) A backup central switch (for SF ITMS, it will be located at the Office of Emergency Services -OES) with 1Gbps links.

With proper bandwidth provisioning in the middle and edge switches, the traffic on any single uplink is designed so that the total amount of traffic cannot exceed its capacity, although it can approach the full bandwidth. Depending on the phased deployment of the backup central switch at OES, two alternative uplink paths from individual middle switches to the TMC may be installed initially for a providing an additional level of redundancy.

The TMC central switch is connected with multiple 1Gbps trunked links to the backup central switch but not enough links to support the full burden of traffic throughout the City to the backup central switch and *then* to the TMC (when fully deployed in San Francisco, traffic could approach 24Gbps to the TMC). This is a reasonable decision since:

- (1) Such failure conditions requiring full routing of all cabinets within the city to OES and then to the TMC cannot be reasonably expected to be possible with the conduit layout of the city, and
- (2) The single entrance to the TMC building will preclude such possibility.

Edge Switch

Cameras, CMS, detection stations, and other devices are connected to the communications systems at traffic signal controller cabinets. As such, there will be, at most, 1,500 edge switches when the SF ITMS is fully deployed. The traffic controller cabinets are not environmentally controlled and thus the edge switches are field-hardened devices, capable of withstanding extended temperatures (up to 70°C or 130°F), humidity, and some vibration and shock tolerance above what would be typical office-grade switching equipment. The edge switches have at least six full-duplex 10Mbps, RJ-45 ports and two

100Mbps ports and sufficient backplane bandwidth for non-blocking, wire-speed switching of all ports simultaneously (at least all 10Mb inputs are supported at full data rate on the 100Mb uplink port). Preferably, the edge switch has built-in single-mode fiber transceivers, operable up to a 15km distance. The edge switches are not required to have Layer 2 management functions initially (e.g., Spanning Tree) but can reasonably be assumed to support such management functions in the near future at the same price point. The devices used during the architecture testing were Garrettcom P62F, which have six 10/100Mbps RJ45 ports and two single-mode fiber, 100Mbps uplink ports rated for transmission up to 20km.

Middle Switch

Throughout the City, there is a middle switch for roughly every *twenty* field cabinets. Up to *sixty* middle switches will be required throughout the City when fully completed. All middle switches will be housed in an environmentally-controlled cabinet (“communications hub”) since currently available Ethernet hardware of this capability is not yet field-hardened. The middle switches that were used during the architecture test were Extreme Networks Alpine 3804 with 48 100Mbps ports with built-in single-mode fiber transceivers rated for 20km+ transmission distances, two, 1GBps single-mode fiber uplink ports rated for 15km+ transmission distances, and one 1GBps single-mode fiber port used for the link to the paired middle switch.

Twenty field cabinets use this middle switch as their “primary” path for transmitting data to the TMC. Twenty ports are reserved as the “backup” path from other field cabinets. Each middle switch has a full non-blocking wire-speed backplane to support 240MBps data from primary links, 240MB of data from backup links, and 480Mbps data from the backup link (for a total of 960 Mbps maximum total traffic uplink on the 1GBps uplink) simultaneously onto the 1Gbps uplink to the TMC. The 240Mbps loading assumption assumes that 20 cabinets transmit 12Mbps of data each (four cameras at 3Mbps each or a lesser number of cameras at higher data rates). The primary role of the middle switches are for data aggregation and bulk re-routing under failure scenarios.

The middle switches support all Layer 2 standards and management functions including Spanning Tree, Class of Service / Quality of Service with eight priority queues, and multicasting. As the network grows in size, the middle switches are upgradeable to Layer 3 functionality.

Central Switch

The central switch at the TMC (and the backup duplicate at OES) is an enterprise-class switch with blades for at least 64, 1Gb single mode fiber ports, eight multi-mode/copper, 1Gb ports (for distribution of data within the TMC and intranet) and upgradeability to 10GB ports when the technology is available (e.g. Extreme Networks Black Diamond 6816). The central switch is essentially a chassis-design so that the switch can be expanded easily as the system grows. Most of the 1Gb ports are used for input from the middle switches (approximately 60) and several are used for data transfer to the TMC video wall

server, cross-links from the alternate central switch, TMC workstations, DPT LAN, and video servers for internet and intranet delivery of streaming video and traffic information.

The non-blocking central switch backplane of 256Gbps is more than capable of handling all incoming data and video all of the time. Considering a maximum of 12Mbps from each edge switch, (1,200 switches; four 3MB MPEG2 videos from each cabinet) the full load on the core backplane is 24Gbps from the field. The central switch used during the architecture test was an Extreme Networks Summit 7i (for the small simulated network, “Black Diamond”-class performance was not required).

Failure tolerance

Failures to be tolerated include cuts to the fiber runs and failure of switches. Cuts to a fiber cable will eliminate a particular path or paths from a number of devices to the TMC. Based on the location of the fiber cut, termination of a communications path could include as few as one device or as many as hundreds of devices. In general, the closer to the TMC the fiber cut occurs, the more devices will be affected. Similarly, failure of switches will affect as few as a single device (failure in a device cabinet) or possibly the entire system (failure in the TMC – although an alternative workstation could easily be used at the backup TMC in OES), based on the location of the switch.

With the logical architecture presented in this paper, at full build-out there are *six* alternative (logical) routes from a particular field cabinet to the TMC that utilize *ten* separate physical links (which are geographically separated as much as possible in the physical conduit and fiber splice design). Thus, communications to a particular field cabinet is only lost during a *combination* of failures. At fewest, failure of both uplinks or both middle switches connected to a particular cabinet will cause a communications loss. At most, a combination of *five simultaneous link outages* is required to disallow communications to a particular field device.

The redundancy provided is a trade-off between network complexity and over-subscription of uplinks during extreme failure scenarios. The single point of failure occurs at the entrance to the TMC where all routes converge into the same conduit. A cut there would be catastrophic to the availability of data. However, at virtually any other point in the network, when the network is fully deployed, a single cut will have *no effect*. When a fiber is cut or a switch fails, the connectivity is restored automatically by the switch network using standard Spanning Tree technology.

When a link is terminated (e.g. fiber cut), the switch will “find out” that the link is down and compute a new best path (revise the spanning tree). These calculations typically occur in less than one minute, but in networks of appreciable size and complexity such as SF ITMS when fully deployed, the re-calculation of the spanning tree could take five minutes or more. During that time, communications between all devices in the sub-domain are not available.

To calculate the new spanning tree, the ports of each affected switch must be shut down. Once a new path is found, communications are restored. For the SF ITMS, such latency should be considered an acceptable interruption because it will only occur when a fiber is cut or a switch fails. Faster operation of spanning tree calculations is available currently through proprietary vendor technologies (such as Extreme Networks' ESRP) that can reduce the spanning tree recalculation time to seconds, and perhaps milliseconds. These technologies will migrate to industry-accepted standards in the near future.

Preservation of mission-critical data streams

The same switches are used for data transmissions (e.g. signal controller status) as for streaming video. The total bandwidth required for data is miniscule compared with the requirements for streaming video and can be easily accommodated on any one remaining path to the TMC in the event of multiple trunk failures. It is very important that such data streams are not neglected in preference for video traffic. Under all failure scenarios, this architecture provides sufficient bandwidth for all cameras to transmit streaming MPEG-2 video to the central switch under multiple trunk failure conditions. No situation is created where a particular link could become oversubscribed as long as the assumptions are met – primarily that only four 3MBps streams are sent from each traffic signal controller cabinet back to the central switch. Using Ethernet, it is easily designed into the network the capacity for:

- (1) A subset of video streams is maintained at full capacity without interruption,
- (2) The mission-critical data transmissions from field devices are preserved, and
- (3) The remaining video streams are dropped entirely (or allowed to be degraded beyond usability).

Hence, those video transmissions that are prioritized to get through have *no* degradation (since minor degradation, loss of a few packets, at MPEG-2 quality results in un-useable video due to the structure of MPEG2 compression). The data transmissions are protected by using *Class of Service* Ethernet standards.

Class of Service is an Ethernet standard that allows priority treatment of particular transmissions streams at switches by creating a number of *queues* internal to the switch. Most enterprise-class switch manufacturers of the quality required by the SF ITMS program use eight queues for traffic prioritization. When a particular packet arrives at the switch, if it is priority 8, it is sent immediately to the head of the line (unless the link is over-subscribed and the switch has a non-blocking backplane, there is never a line). Using this standard switch technology, data transmissions from traffic controllers, CMS signs, and other health and status low-bandwidth communications can be sent to the front of the queue at all times so that those transmissions are not “overwhelmed” by the lower-priority video streams, which are approximately 300 times the size of data transmissions. Class of Service priorities can be assigned by MAC address, IP address, port, and a number of other methods – making it a very flexible way to guarantee particular data transmissions are maintained even under extreme failure scenarios (e.g. enabling such rules as, “under failure

scenario A, allow all videos from PTZ cameras to continue transmission and drop all videos from fixed cameras” or “under failure scenario B, allow all videos from IP address subset 128.156.XXX.101 and drop all others”, etc.).

Extended temperature and humidity tolerance

All Ethernet middle and central switches are sensitive electronics and are designed to operate in controlled environmental conditions. Since the electronics generate heat, the primary problem is over-heating due to extended ambient temperature. Placement of Ethernet equipment in the field will require some form of hardening for extreme heat or air conditioning of the enclosure (cabinet and/or field hubs) to reduce the ambient temperature. Moisture (humidity) is not typically a problem for equipment until the humidity is high and the ambient temperature is higher than that of the device itself (i.e. condensing conditions). Although the City of San Francisco does not routinely experience extreme temperatures, it does at times reach 100°F several days during the year. In addition, metal cabinets exposed to direct sunlight can be heated to significant temperatures (including the effect of the cabinet’s typical dark green paint color). Deployment of Ethernet communications architecture for SF ITMS field switch deployment will require environmentally controlled cabinets with air conditioning devices and circulation fans.

Demonstration Test Bed

In order to illustrate that the proposed logical and physical communications architecture is feasible and operates as intended, a demonstration test bed was set up in conjunction with:

- (a) Switch vendors (Garretcom and Extreme Networks),
- (b) CCTV camera vendors (Pelco, Kalatel, Cohu, and Philips),
- (c) Video server vendors (TLC Watch, Broadware),
- (d) Traffic management software vendor (Siemens – GTS),
- (e) Codec vendors (Vbrick, Cornet, Enerdyne, TLC Watch), and
- (f) Traffic controller vendors (SafeTran and Econolite)

Test Suite

Figure 2 illustrates the test suite that was installed at the Extreme Networks Interoperability Laboratory in Santa Clara, CA for five days. The test suite included:

- Four CCTV cameras with Pan, tilt, zoom functionality.
- Two fixed CCTV cameras.
- Six codec pairs using H.232, MPEG2, MPEG1, and MJPEG formats.
- Two 2070 traffic controllers with Ethernet-based communications traffic control software (NextPhase version 1.4.3).
- Two Garretcom P62F field-hardened “dumb” switches with dual single-mode fiber transceiver outputs, six 10/100MB RJ-45 inputs, and non-blocking backplane.
- Two Extreme Networks Alpine 3804 intelligent “middle” switches operating at Layer 2 with QoS, Spanning Tree, Trunking, IGMP, SNMP standards. Full

wire-speed non-blocking backplane capacity with 24 100Mb single-mode fiber ports, four 1GB fiber uplink ports.

- One Extreme Networks Summit 7i intelligent central switch operating at Layer 2 with QoS, Spanning Tree, Trunking, IGMP, SNMP standards with full wire-speed non-blocking backplane capacity. Operating at all Gigabit input/output ports (surrogate for the Black Diamond-class central switch).
- One Extreme Networks Summit 1i intelligent “middle” switch (surrogate for the lateral link between Alpine 3804 middle switches which would occur in the field) operating at Layer 2 with QoS, Spanning Tree, Trunking, IGMP, SNMP standards with full wire-speed non-blocking backplane capacity. Operating at all Gigabit input/output ports.
- Four NTSC monitors.
- Two programmable load generators suitable for simulating streaming video and burst data traffic on the middle switch and central switch.
- One PC (Dell Latitude Cpx) with Icons traffic management software and video viewer software controller/client software with a 100Mb NIC.
- Two Cat-5 to Single-mode fiber 10MB Ethernet media converters with appropriate connectors (for simulating traffic controller cabinets without video feeds).




Figure 2. Traffic controllers, switches, and laptop clients

A number of tests were executed to verify that the communications architecture would operate as designed. The tests were primarily focused on verifying that the architecture operates satisfactorily (i.e. data transmissions and selected video transmissions are protected) under extreme failure conditions where the uplink from a particular middle switch is saturated. For brevity, the details of test execution will not be provided here. All tests for architecture viability were satisfactory. More details on the test results can be

obtained from DKS Associates, Siemens-GTS, or the City of San Francisco Department of Parking and Traffic.

Summary

An innovative, end-to-end Ethernet-based communications architecture was designed for highly reliable, fault-tolerant communications between a central transportation management center (TMC) and field devices (traffic signal controllers, CCTV cameras, changeable message signs, etc.). A field-hardened ethernet switch is placed in each cabinet with video cameras and connected to two middle switches via single mode fiber optic cable. Each distribution “middle” switch supports 20 field cabinets as a “primary” uplink to the TMC, 20 backup connections, and a “lateral” backup from another middle switch (an additional 40 cabinets) assuming four 3MBps video streams per intersection. Each middle switch is connected to the TMC central switch and a backup central switch at an alternate location. Six logically separate uplink paths utilizing 10 physically separate links (as separate as possible given the conduit design) in the network are provided from each signal cabinet for a high degree of redundancy and communications reliability. Automated fault tolerance and data transmission preservation is provided by industry-standard Quality of Service and Spanning Tree Ethernet protocols. The architecture was tested and demonstrated at the Extreme Networks Interoperability Laboratory in Santa Clara, CA and was found to be a viable alternative to ATM, SONET, frame relay, and other networking technologies for integrated transportation management applications.