the Availability Digest

QEI Provides Active/Active SCADA with OpenVMS

September 2007

For almost 50 years, QEI of Springfield, New Jersey, (www.qeiinc.com) has been supplying highly available SCADA systems to the electric, transit, gas, water, and other utilities. A SCADA (Supervisory Control and Data Acquisition) system provides controllers with the facilities required to monitor and control the field devices upon which these utilities depend. It automatically generates alarms should conditions in the field demand immediate controller attention and provides a raft of historical data for trend analysis, root cause analysis, and many other functions important to the utilities.

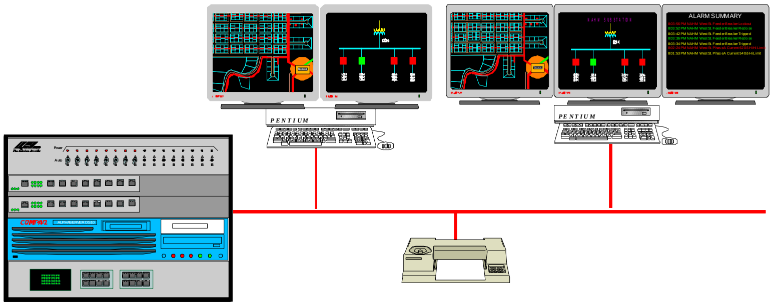

QEI's current SCADA system, TDMS-PLUS (Total Distribution Management System), focuses on the monitoring and control of electrical power substations used for the distribution of power to electric utility customers and transit systems. Built on the highly reliable and secure HP OpenVMS platform, TDMS-PLUS provides extreme availabilities through the use of dual, triple, or quadruple active/active redundancy in disaster-tolerant configurations.

We first describe the TDMS application and its hardware and software architecture. We then focus on its high-availability features.

The Electrical Distribution Network

The household and building electrical power upon which we all depend comes to us through three major infrastructures – generation, transmission, and distribution:

- Electricity is *generated* by large electrical generators which are themselves powered by coal, oil, gas, water, or other sources.

- This electricity is carried over long distance *transmission* lines to local points of distribution. Since power loss over the transmission network is a function of the current flowing through the lines, transmission networks distribute electricity at very high voltages and low currents (power is voltage times current). Typical transmission voltages are 230 to 500 kilovolts.

Transmission networks terminate in transmission substations, which reduce the voltage for distribution to homes and businesses.

- From the transmission substations, the lower voltage electricity is distributed to other distribution substations for routing to homes and offices. These substations feed the power lines so ubiquitous on the telephone poles outside of our homes. The fact that these lines carry electricity at a “lower voltage” is a bit misleading. The voltage on the power lines outside of your home is 13,000 volts. Don’t touch a downed line!

  This voltage is reduced by transformers on the poles near our homes or businesses to the 120/240 volts that we expect.

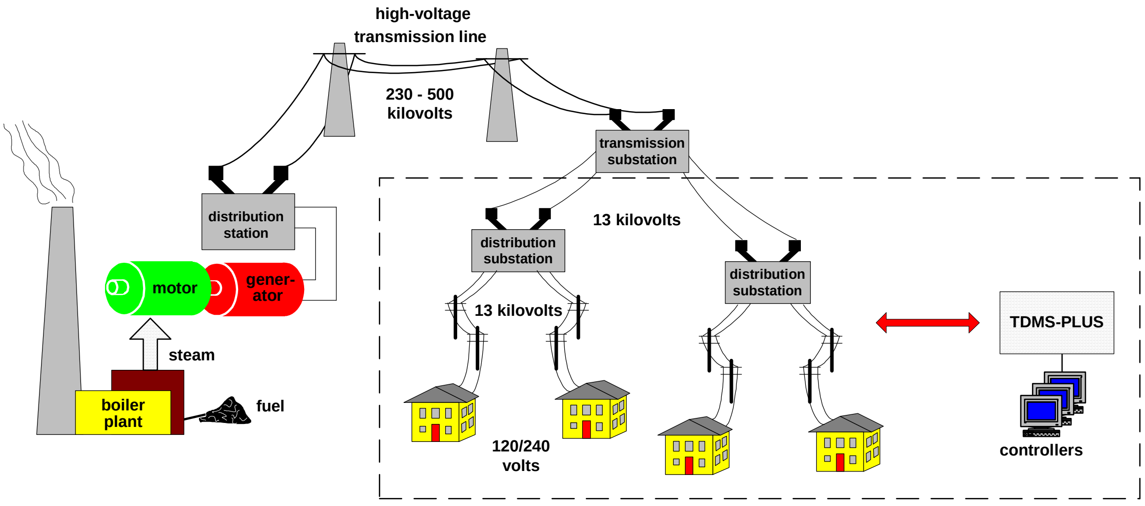

Electrical Distribution System

QEI’s TDMS-PLUS SCADA systems are used to monitor and control the electrical network from the distribution substations to our homes and businesses. They are also found in many major metropolitan transit systems as well as long-haul train systems performing the same function for power distribution to the trains and subways.

QEI’s TDMS-PLUS

The QEI TDMS-PLUS SCADA system comprises three major hardware components:

- The Master Station, which provides the intelligence for the system.
- The Worldview Controller Consoles, through which the controllers monitor and control the system.
- The Remote Terminal Units, or RTUs, which monitor the status of digital and analog indications in the field, send field changes to the Master Station, and receive and execute commands on field devices from the Master Station.

Master Station

The Master Station runs on HP OpenVMS blades or towers. It can be configured as a standalone system or in a dual, triple, or quadruple modular redundancy configuration, as described later.

The Master Station polls the RTUs for status changes and maintains displays of the current field status on the Worldview terminals for the controllers. It accepts commands from the controllers or from monitoring applications to send to the field devices.

Each field device is referred to as a *point*. There are three types of points:

- Digital points, such as circuit breakers, are in one of two states (though multistate points can be configured as a set of dual state points). Whenever the state of a digital point changes, its new value is sent by the RTU to the Master Station.

- Analog points, such as voltage or temperature measuring devices, report a measured value. Specified for each analog point is a dead zone within which the value of the analog point measurement can properly lie. Should the value of the analog point move outside of the dead zone, the analog point value is sent to the Master Station.

- Computed points are pseudo-points established within the Master Station logic. The value of a computed point is calculated based on the status and value of other points but is otherwise treated as if it were in the field.

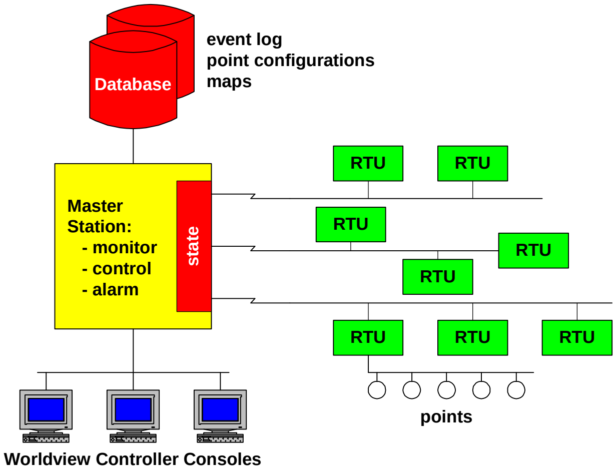

The Master Station maintains in memory the status of every point. It continually monitors the digital, analog, and computed points for alarm conditions. Should it detect an alarm condition, it will notify the appropriate controller by means of an audible alarm and a change in the displayed status of the offending point (for instance, by changing its color and causing it to blink).

The controller can attempt to correct the alarm condition by sending commands to pertinent field devices. For instance, if the power drain on a particular distribution line is becoming too large, he may be able to shut down some of the devices powered by that line to reduce the power drain and return it to an acceptable level.

The Master Station also maintains all persistent information on mirrored disks. This information includes a log of all events (state changes and operator actions) with time stamps, the current configuration of all of the points, and the maps used by the Worldview consoles. Also included are tags. A tagged device cannot be energized mistakenly by a controller. This prevents maintenance people in the field from being injured. The current system state is not maintained on disk since it is highly transient.

The database may be replicated in real time to commercially available databases such as Oracle or SQL Server for such offline analyses as trend analysis and root cause analysis of faults. In addition, the event log can be played back on the Master Station to show exactly what happened during an incident.

Worldview Consoles

The Worldview consoles provide graphic displays of the electrical network and tabular displays of detailed information. Graphic displays include maps and “one-line” diagrams. The maps are used to display the complete or partial electrical network and show where the trouble may be. Substations and other devices are shown by symbols. One-line diagrams are used to display devices within a substation and show their interconnection. This is the most commonly used display.

Some of the field devices, such as circuit breakers, are controllable. Via the console, the controller can send commands to these devices to change state. Each command follows a select-before-operate (SBO) sequence. The device is first selected and reports back to the Master Station that it has been selected. The selected device is shown on the controller's console so that he can be assured that he is controlling the proper device. He can then execute the command and view the resulting change in state on his console.

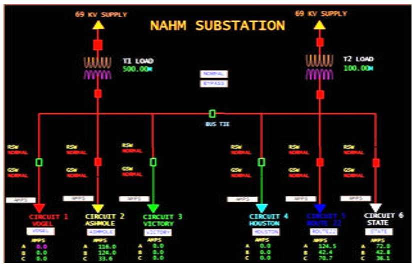

The Worldview consoles are also used to configure points (name, location, analog or digital, alarm state, display characteristics, and so on). Maps and one-line diagrams may be built via a Worldview utility, or they may be imported from other packages such as AutoCAD. The maps and one-line diagrams can be panned and zoomed and are layered to provide decluttering as a controller zooms out.

The maps can be divided into zones. An operator can only control devices in the zone to which he is assigned, though he can view device states in other zones.

Worldview runs on standard Windows PC terminals.

Remote Terminal Units

The remote terminal units sit in the field and monitor the status of digital and analog points. They may be contained in large cabinets positioned in a substation, or they may be small pole-mounted units monitoring a transformer that is stepping down voltage to a home.

An RTU monitors multiple digital and analog points. A communication line from the Master Station may connect to only one RTU if that RTU is monitoring critical points or to dozens of RTUs in rural areas.

All of the RTUs on a single communication line are polled repetitively for changes. When polled, an RTU will send status changes, if any, to the Master Station. It will also execute commands sent from the Master Station by selecting the specified point, confirming the selection to the Master Station, and then executing the command upon the receipt of an execution message from the Master Station.

A typical TDMS-PLUS SCADA system may have dozens of RTUs monitoring thousands of points.

Communication between the Master Station and the RTUs can be via any communication medium, such as telephone line modems, LANs, WANs, fiber, microwave, or radio. Communication over modems occurs at bit rates from 1,200 bps to 56 kbps. If modem connections are used, these relatively slow communication rates limit the SCADA system event rate to typically twenty to fifty events per second. However, LAN speeds of 10 to 100 mbps eliminate this limitation.

RTUs support management from the Master Station via SNMP (Simple Network Management Protocol) over LAN and WAN connections.

RTUs are not redundant, though they support redundant communication channels. Even in the extreme heat, cold, and stormy conditions in which RTUs must operate, QEI data accumulated over 25 years shows an RTU MTBF (mean time before failure) of 200,000 hours (about 23 years). Most failures are caused by lightning strikes.

TDMS Software Architecture

TDMS runs under the OpenVMS operating system. OpenVMS is a very mature operating system that is tailored for real-time operation and is extremely secure, an attribute demanded by power utilities.[1] OpenVMS runs 65% of the world's energy management systems.

The TDMS software provides data acquisition, point control, alarm processing, tagging, zoning, event logging, and reporting, among other functions. Add-on applications implement customized functions such as load management through voltage adjustments and/or through the disabling of customer equipment.

Through its software services, TDMS maintains an image of the state of all field devices in its memory. As each point state change is received from the RTUs, TDMS' memory state map is updated. It is also updated with the results of any computed points. Periodically, the state map is verified via an all-data poll, which asks all RTUs to send the current state for all of the points that they are monitoring to the Master Station.

Changes to device state are written in an event log to disk. In addition, changes to point configuration maps, and one-line diagrams, are recorded on disk. The memory-resident state map is not maintained on disk since it changes rapidly and is transient in nature. Furthermore, it is easily recovered in seconds if necessary via an all-data poll.

Active/Active Redundancy

TDMS also supports active/active redundancy by configuring a second system synchronized with the active system.

Active/Active Dual Modular Redundancy

A second TDMS Master Station can be configured to back up the active primary Master Station. All of this system's processes are enabled, and its database is synchronized with the active system's database. The system is ready to take over processing instantly.

However, in a SCADA application, this second system cannot be actively processing events along with the primary system. This is because all events must be processed in sequence, just as they are generated in the field. Otherwise, a false condition may be reported. Therefore, only one Master Station can be active at any one time. Though this is truly an active/active architecture, it must necessarily be used as an active/standby configuration with nearly instant takeover. We will therefore refer to the second system as the standby system.

The active and standby systems may be collocated, or they may be geographically separated for disaster tolerance purposes. The systems are connected by LAN if they are collocated. They are connected by WAN, such as a T1 link, if they are geographically separated.

The controller Worldview consoles have access to both systems via a LAN, which may be either a dedicated or a corporate LAN. Alternatively, if the standby system is remote, a separate set of consoles may be provided at the remote facility.

[1] OpenVMS is the most secure operating system, according to the Department of Defense's Computer Emergency Response Team (CERT). There has never been a reported incident of an OpenVMS system being infected with a virus.

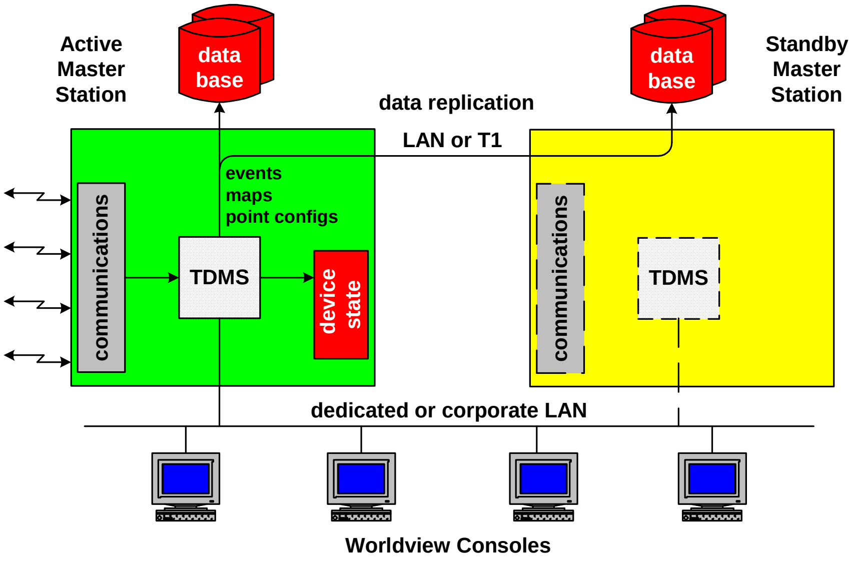

Database Synchronization

The TDMS database contains, among other things, a log of all events, the current point configurations, and the current maps and one-line diagrams. As events, point configuration changes, map changes, or one-line diagram changes are processed by the active Master Station, they are written to disk. They are also sent over the communication channel connecting the active system to its standby. At the standby system, this same data is written to its disk so that the databases are in synchronization. Should the active system fail, no events or configuration changes will be lost (except, perhaps, those that were in the replication pipeline at the time of failure).

The memory-resident state maps are not replicated since they are highly transient in nature and are easily reconstructed following failover, as described next.

Failover

The active and standby systems also trade heartbeats over their interconnecting channel. Should the standby system detect a loss of heartbeat from the active system, it takes over control of the system. Its processes are all enabled, and it has an up-to-date database. However, it does not have the current state of the field devices.

Therefore, it must issue an all-data poll to its RTUs. It polls its RTUs for current device state simultaneously on all communication channels. Depending upon the number of points being monitored by the system and the speed of the communication lines, this typically takes anywhere from thirty seconds to two minutes.[2]

As it receives device status, the new Master Station builds its own in-memory state map and makes device state available for showing on the system maps.

Once all updates have been received from the field, the system is fully operational.

[2] There is currently a move in the industry toward higher speed lines such as fiber. This will result in much faster recovery times.

Modular Redundancy

The TDMS SCADA system can be configured with dual modular redundancy, as described above. It also can be configured with triple or quadruple modular redundancy. A typical triply redundant system would include an active/standby pair at the primary site and a standby system at the remote site.

A system with quadruple redundancy would comprise an active/standby pair at the primary site and an equivalent standby/standby pair at the remote site.

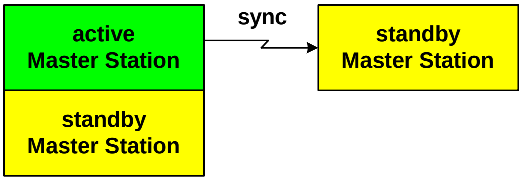

triple modular redundancy

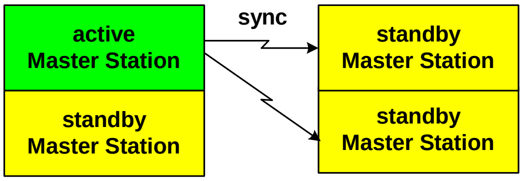

quadruple modular redundancy

As indicated above, no matter what the level of redundancy, only one master station may be active at any one time.

Availability Experience

The TDMS-PLUS OpenVMS systems are maintained by HP. If HP's 24x7 maintenance is used, HP guarantees a five 9s (five minutes per year) availability for its hardware and software through its Service Level Agreement. An 8-to-5 next-business-day maintenance agreement is also available.

In all of its experience, QEI reports only one redundancy failure attributable to TDMS. That was when two systems tried to take over the active role. Other failures are quite infrequent, and are generally due to inappropriate configuration changes made by the customer.

Summary

QEI has been delivering SCADA equipment for almost 50 years, beginning with tone equipment for data transmission under its original name, Quindar.

Since then, it has gone through many evolutions of the Master Station:

- 1965 – Hardwired control panel.
- 1968 – Integrated circuit-based control panels with computer loggers.
- 1974 – First minicomputer-based Master Station using Texas Instruments 960.
- 1981 – Master Station using DEC PDP-11.
- 1986 – Master Station using DEC VAX and VMS.
- 1993 – Master Station using DEC Alpha and VMS (QUICS-IV).
- 2000 – Master Station using DEC Alpha and OpenVMS (TDMS-2000).
- 2004 – Master Station using HP Integrity and Open VMS (TDMS-PLUS).

QEI's TDMS system reflects decades of field experience. Its current systems are aimed at monitoring distribution subsystems for electric utilities and transit systems. With redundancy levels up to quadruple modular redundancy and with fast failover, QEI's customers are experiencing extremely high availabilities in actual practice.