

Advanced Diagnostic/Prognostic Solutions for Complex Information Technology (IT) Networks

As computer networks have become an increasingly important element of business and industrial applications, customers are becoming increasingly concerned with data and network reliability and availability. In early 2012, the Tokyo Stock exchange was affected by a failure of a key IT component. Due to this problem, the exchange was closed for 2.5 hours.¹ The cost of that outage was enormous in lost trading volume and productivity, and unfortunately similar instances have occurred elsewhere.

Maintaining “Uptime” on IT Networks

IT networks have also moved beyond the office to the factory floor. Networked manufacturing cells under central control govern the production levels. Untimely failures of network components can cause tremendous loss. Clearly, reliability and robustness of networks, and the scheduling of maintenance and upgrades of those components, is absolutely critical. Forward-thinking firms are embracing the notion of predictive maintenance schedules based on the monitored condition of the individual assets on the network. This strategy is referred to as condition-based maintenance or CBM, and it can reduce maintenance costs by focusing attention on the identification and repair or replacement of degrading assets.

Maintaining a reliable IT network is challenging, given the wide variety of diverse hardware (blade servers, routers, switches, UPS power systems, etc.) and the software components that are supported by these IT networks. Additionally, networks must be especially robust while operating in an environment where expert support teams may not always be available to troubleshoot and repair a network malfunction or an anomaly that may cause degradation in system performance.

What is needed to address these challenges is an automated tool that continuously monitors the network, assesses changes that affect performance, and reports problems via cell phone text messages or email to the IT network administrator for corrective action.

This application note describes an innovative web application called Sentinel Network™ that provides the customers with a powerful “net-centric” prognostics and health management (PHM) software platform solution.

The Sentinel Network tool performs:

- Accurate network discovery
- Network configuration monitoring
- Real-time resource health monitoring of workstations/servers
- Uninterrupted power supply (UPS) monitoring
- Network health monitoring based on load changes
- Background monitoring and data collection
- Switch troubleshooting

¹ Imagami, K. “Server Problem Caused Tokyo Glitch,” Wall Street Journal, February 8, 2012

Technical Approach

The foundation of Sentinel Network rests upon an extensible software platform that distributes the sensor data collection, data fusion, reasoning, and presentation tasks. The tool also supports timely introduction and efficient software maintenance through its scalable design, enabling parallel development and support. Thus changes or the addition of more assets to the network are automatically supported. The distributed software architecture addresses these features with a modular web server design that separates the main client application from reasoner extensions and services that manage sensor data collection.

Sensor data is collected by leveraging the inherent simple network management protocol (SNMP) transport from the network's workstations, servers, switches, and UPS devices. The acquired sensor data is collected and stored to a central PHM database in Sentinel Network. The data is then processed using advanced diagnostic and prognostic reasoners, which process the multivariate sensor data to isolate the root cause of the fault condition and then estimate the remaining service life of the device being monitored.

For network devices, including UPS systems, this results in a remaining useful life (RUL) time estimate for impending failure and allows a technician to make the proper repairs or replacement to avoid downtime from equipment failure. The Sentinel Network PHM solution presented can be used to support a variety of sensor networks supporting critical systems, such as factory-floor automation cells, aircraft, automotive, and many other complex systems that cannot afford unscheduled maintenance due to system or subsystem malfunctions.

Critical networks can include routers, switches, servers, and potentially several thousand client nodes of varying type and function. The mission-critical components of the network are powered through UPSs. This device is responsible for providing clean, uninterrupted power for the connected network devices even during a power outage. The UPS is known to be a single point of failure for an IT network, as it powers the core of the network responsible for mission-critical systems. UPS devices have failure points similar to other power systems including the batteries, capacitors, MOSFETs, and DC/AC inverter stages. The ability to predict the RUL of a UPS will improve the operational availability (Ao) and reliability of the network by reducing unplanned network downtime.

The foundation of this comprehensive PHM solution for the IT network systems is predicated upon an extensible software platform that distributes the sensor data collection, data fusion, reasoning, and presentation tasks. Timely introduction and efficient software maintenance is contingent upon a scalable design that enables parallel development and support. The distributed software architecture addresses these objectives with a modular web server design that clearly separates the main client application from reasoner extensions and services that manage sensor data collection.

Figure 1 shows a basic overview of a Sentinel Network implementation for a factory floor, monitoring both IT and operating equipment such as robots. The robots can incorporate complex actuators, power systems, and electronic controls, all of which can be monitored with Sentinel Network. The screen shot in Figure 1 shows the main view of Sentinel Network. A discovered network of routers, switches, servers, workstations, and UPSs displays in the left pane, the real-time monitoring of those devices in the right pane, and the alerts generated as a result of monitoring the network in the bottom pane.

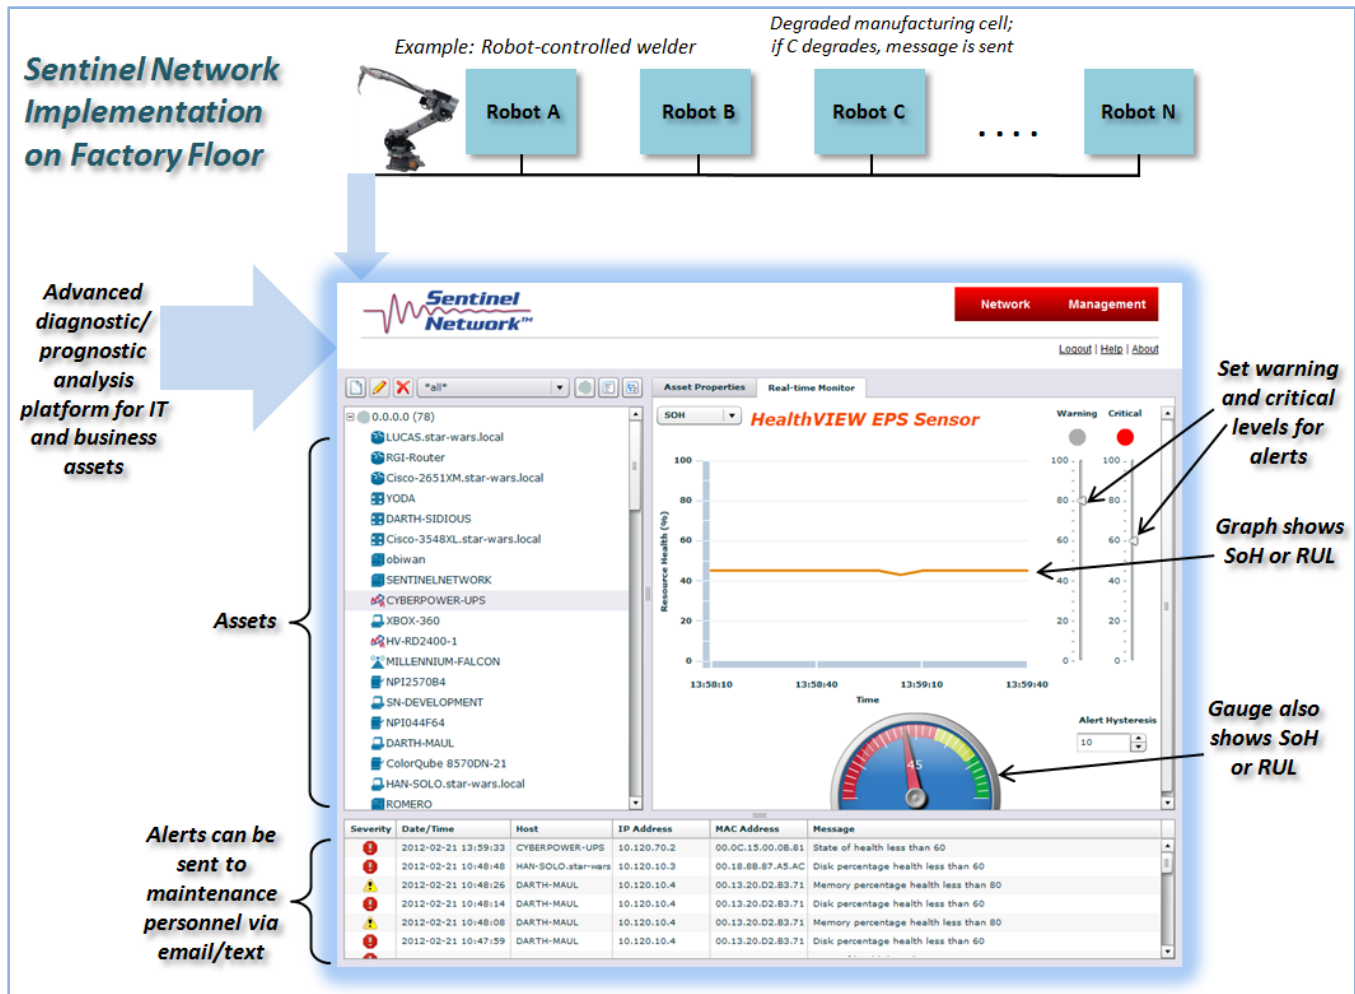


Figure 1: Sentinel Network screen showing the discovered assets and the HealthVIEW indication of network health; manufacturing assets can also be monitored

For example, sensor data is collected using an SNMP transport from workstations, servers, and UPS devices. Those devices that do not natively support SNMP will require the addition of a simple SNMP gateway to publish the data from the devices into a management information base (MIB) that Sentinel Network can then access. The collected sensor data is stored to a central PHM database. The data are then processed using a suite of sophisticated diagnostic and prognostic reasoners, installed on the Sentinel Network host, that process the multivariate sensor data to isolate the root cause of the fault condition and estimate the remaining service life of the device being monitored. For various devices, this results in a clear prioritization of maintenance for the technician so the network avoids catastrophic downtime. The technician can then make the proper repairs or replacements to sustain the network and avoid downtime.

A typical IT network rack with the UPS included is depicted in Figure 2. This device is being monitored in real time for RUL prediction using Ridgetop’s Adaptive Remaining Useful Life Estimator (ARULE™)² algorithm.

Sentinel Network features are described in Figure 3.

² James Hofmeister and Sonia Vohnout, “Adaptive Remaining Useful Life Estimator,” ISHM 2009




Figure 2: Typical server rack

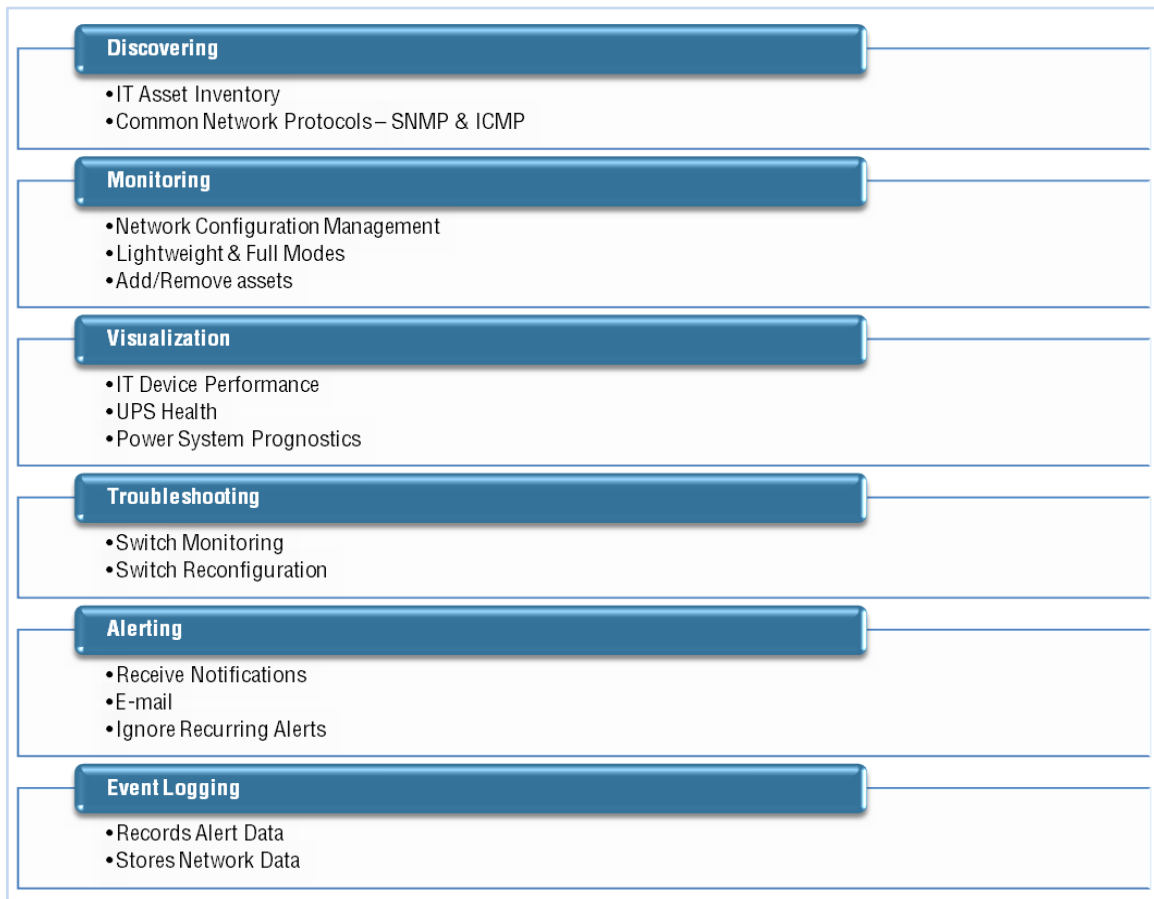


Figure 3: Sentinel Network features

Currently, the competitive advantage of the PHM solution presented in this application note is summarized by its system prognostics and troubleshooting capabilities.

First, Sentinel Network has the capability of incorporating prognostic reasoners for network devices through advanced implementations on UPS and power systems. The reasoners use algorithms that are based on physics-of-failure (PoF) models, rather than older data trending-based PHM approaches. The data-trending approach to determine the health of a device is accomplished by considering historical data and its trends, while PoF considers the device life-cycle load and its performance to identify potential failure. Physical devices can be monitored in real time through PoF models. Device data are collected through SNMP and are processed by a prognostic reasoning engine.

Switch Monitoring and Troubleshooting

Sentinel Network's troubleshooting capability makes this web application a proactive network monitoring tool. Nearly all managed devices have a file structure that contains both a working and certified directory structure. Alternatively, some managed devices use the reference running and startup configuration files. Managed devices compare the working and certified boot configurations. If they are the same, the device boots from the working configuration. If they are different, the device boots from the certified configuration. It is possible to boot from the working configuration if the configurations are different, however this has to be forced from the managed device. Recovery from a serious network failure such as changes to the switch configuration can be accomplished by understanding the managed switch directory structure and how to load either the working or certified boot configuration.

During the initial network device discovery, managed switches such as Alcatel Lucent's OmniSwitch 6850-24 can be added to the switch monitoring feature within Sentinel Network. Once this device is committed to Sentinel Network, an initial configuration from the device is taken as the default or baseline configuration from which to refer during switch monitoring. Switch configuration changes are detected when switch monitoring compares the new switch configuration file to the one collected during the initial discovery. When a change is detected, Sentinel Network alerts that the switch configuration has changed. Sentinel Network has the capability to help the network recover from a serious network failure when these configuration changes are detected. Sentinel Network utilizes an auxiliary connection to recover the managed switch through its serial port. This is done by reloading the managed device through the serial port. The user interface provides an automatic recovery option which connects to the managed device through the serial connection and reloads the switch to the certified directory. Alternatively, the user interface provides step-by-step instructions for performing this process manually.

Hardware Implementation

Several hardware solutions were considered that would enable Sentinel Network to be deployed as a plug-in network device. This allows for full control of the target system. It also makes troubleshooting and supporting Sentinel Network long-term much more desirable. A one-rack unit (1RU) server was selected for the following reasons: size, energy consumption, and performance-to-cost ratio.

The Sentinel Network hardware implementation is shown in Figure 4. This 1RU box will mount in standard 19" server racks. The box is powered by a 1.8 GHz dual core Intel Atom processor, which provides sufficient processing power as well as low energy consumption. The baseline memory is 2 GB of DDR3 1333 MHz with an optional upgrade to 4 GB available. A 160 GB 7200 RPM hard drive also comes standard. Finally, a 250 W power supply powers the system. This hardware combination is available as a standard OEM package, making it a reliable system to deploy Sentinel Network. The total cost of this system is also very low.



Figure 4: Sentinel Network server appliance

Results

Sentinel Network has been proven to customers as a viable solution for prognostics on IT networks. The Sentinel Network product provides network discovery, committing of assets, user interface description, lightweight monitoring, real-time Windows management instrumentation (WMI) monitoring, real-time UPS monitoring, and background monitoring (see Figure 5).

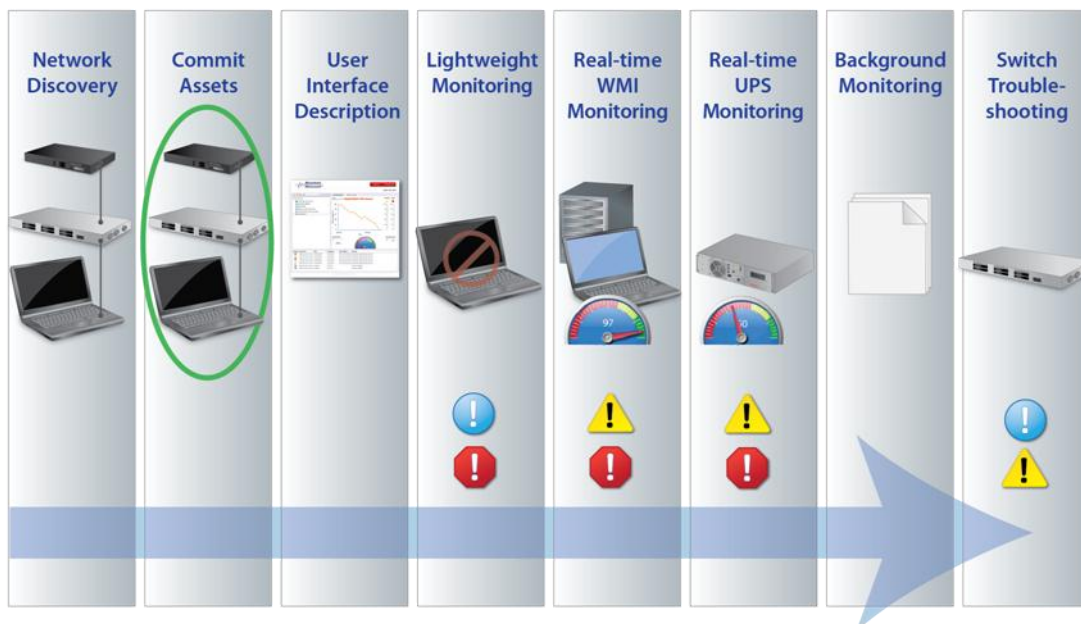


Figure 5: Sentinel network functions

Network topologies that have two separate networks can also be supported, as shown in Figure 6. This dual network was constructed using two routers, three switches, four servers (Sentinel Network is included in this total), four UPSs, and six laptops.

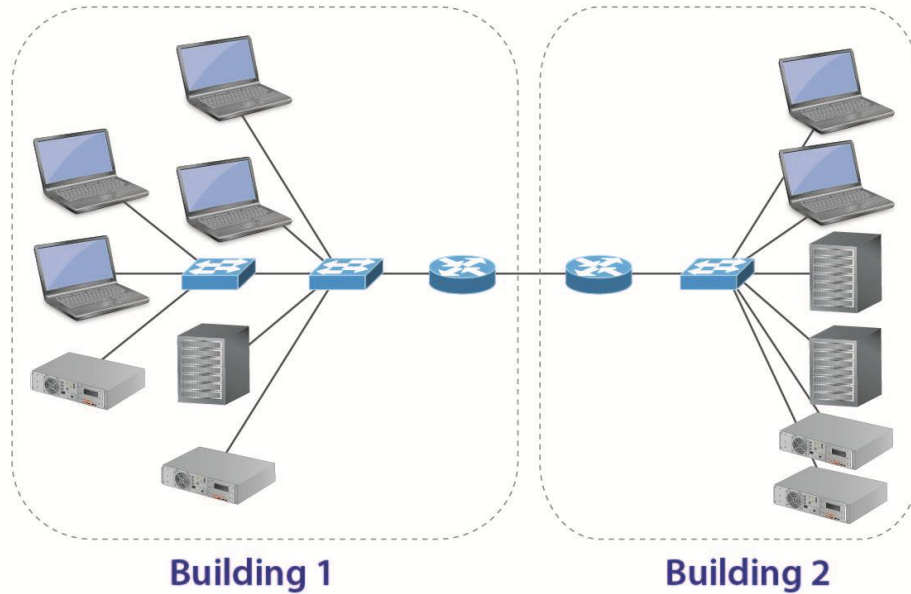


Figure 6: Two-building network topology diagram

00397

Both of the routers for the networks are Cisco 3825 integrated services routers (ISRs). The advanced feature set of this device can support routing for several thousand devices on the local network as well as supporting communication between networks.

Two brands of switches were configured, first a pair of Cisco Catalyst 3750-E switches and an Alcatel-Lucent 6850-48 switch. The Cisco switches were directly connected to the router and provide edge access to laptops, servers, and UPSs. Also used was an Alcatel-Lucent 6850-48 port switch, providing edge access to laptops and UPSs. Figure 7 shows a summary of the equipment used and the building the equipment was associated with.



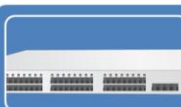

	<p>Cisco 3825 Router</p> <ul style="list-style-type: none"> • Robust integrated services router (ISR) • Provides building 1 to building 2 routing
	<p>Cisco Catalyst 3750-E Switch</p> <ul style="list-style-type: none"> • 24 port, layer 3 switch • Provides building 1 and 2 edge access for laptops, servers, and UPS
	<p>Alcatel-Lucent 6850-48 Switch</p> <ul style="list-style-type: none"> • 48 port, layer 3 switch • Provides building 1 edge access for laptops and UPS
	<p>Dell PowerEdge Servers</p> <ul style="list-style-type: none"> • High performance server, used for file serving, virtualizing applications, etc. • Provides building 1 and building 2 server capabilities

Figure 7: Summary of equipment used and associated building during live demonstration

The design of the Sentinel Network example network was accomplished using 11 subnets, 10 virtual local area networks (VLANs), and class A, B, and C IP addressing. Table 1 lists the VLAN name, id and subnet for each of the buildings. Note that not all 11 subnets are shown in the table. The outside interfaces of each building’s router were configured with class A 192.168.16.0/30 IP addresses.

Table 1: Example Network Analysis Configuration: VLAN Name, ID, and Subnet

VLAN Name	VLAN ID	Subnet
Bldg. 1 Communication Laptops	17	10.100.17.0/24
Bldg. 1 Communication UPS	18	10.100.18.0/24
Bldg. 1 Security Laptops	19	10.100.19.0/24
Bldg. 1 Security Servers	20	10.100.20.0/24
Bldg. 1 Security UPS	21	10.100.21.0/24
Bldg. 1 Telco	25	10.100.25.0/24
Bldg. 2 Science Laptops	46	172.16.46.0/24
Bldg. 2 Science Servers	47	172.16.47.0/24
Bldg. 2 Science UPS	48	172.16.48.0/24
Bldg. 2 Telco	50	172.16.50.0/24

All network devices were configured for SNMP v1 or v2c and the "public" community string. Device host names for laptops, servers, and UPS were defined as <bldg1>-<group>-<device>. An example is BLDG1-SEC-LPT2. Device host names for routers and switches were defined as <bldg1>-<device>. An example is BLDG2-SW1. Laptops were installed with Windows XP Professional and servers were installed with Windows Server 2008. While the total number of devices in each building is small, this network can easily scale to thousands of nodes, making it very representative of actual networked buildings in terms of its design.

Sentinel Network employs two methods for network configuration monitoring. These methods are derived for the purpose of locating new assets on the network, rediscovering committed assets, and determining whether committed assets have gone off the network. During the live demonstration, the network was synthetically degraded by disconnecting a laptop. Lightweight monitoring determines whether the committed assets are still on the network.

Monitoring the resource health for a particular device on the network in real time is critical. Sentinel Network can display the CPU, HDD, and memory as a health value for those measured parameters. During the live demonstration, a laptop and server were monitored for their resource health. Alerts were generated when the chosen parameter dropped below the warning threshold (80%) or the critical threshold (60%). Note that thresholds can be configured by the users depending on their needs.

The real-time monitoring of two UPS devices as part of the two-building network took up the bulk of the live demonstration, and it was divided into state of health (SOH) and RUL estimation based on live data from each UPS.

The first UPS being monitored had a constant load on its output. The SOH calculation was directly related to this value. The RUL estimation is a more complex calculation. The basic assumption here is that the value used as an input to the RUL calculation is indeed the value that is degrading over time. Without this assumption, the analysis fails. The RUL for any UPS starts with the model file consisting of the amplitudes and widths of the three data spaces shown in Figure 8. The RUL is the sum of the widths of the boxes and is recalculated for each new data point. For simplicity, the UPS model was given an initial RUL of two years (in seconds). The widths of the boxes were arranged according to the UPS load under low, medium, and high operation. In other words, the RUL is expected to decrease faster as the load on the UPS increased over a sustained period of time. Only the condition-based data that has amplitude greater than the floor threshold (no faults) and condition-based data that has amplitude less than the ceiling threshold (failure) will change the widths of the data spaces. For each new data point within this region the width of the data space either increases or decreases and resultantly increases or decreases the RUL. With ideal data, there is no correction to the model, however real data causes the model to change, and with each new

data point the changes to the model cause an RUL adjustment. There exists a line on the diagonal of the data space shown in red. When a new data point does not fall exactly on this line, the slope of the line between the origin (lower-left corner of the box) and the data point is calculated. If the slope is within an acceptable tolerance the model is not adjusted. If the slope is outside the bounds provided by the tolerance, the model is adjusted so that the width of the data space either increases or decreases. This change to the data space only occurs when success differences between the model slope and the computed slope are outside the bounds provided by the tolerance. In other words, the model and computation of the RUL require momentum of the data in one direction before making any serious deviations from the present model. Due to the width of each data space in a two-year model, each new data point within a single data space will not greatly affect the overall RUL value. Instead, only jumps between data spaces cause a noticeable change in the RUL value.

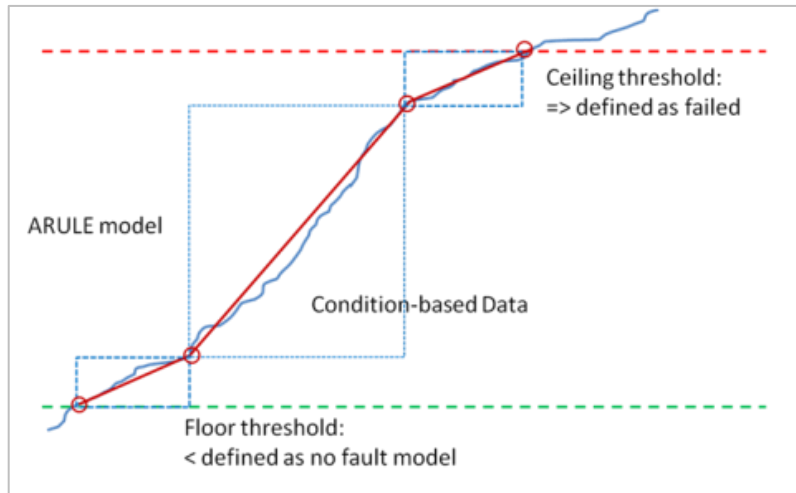


Figure 8: UPS remaining useful life data spaces

A unique default model file for each UPS using the following flat file definition was used.

```

a0 - minimum amplitude for RUL calculation to begin
a1 - amplitude of box 1
a2 - amplitude of box 2
a3 - amplitude of box 3
t1 - width of box 1
t2 - width of box 2
t3 - width of box 3
box - current box
tref - time of first point in a box (i.e., amplitude between a0 and a0+a1+a2+a3
LOOP - value used to determine whether the first point in a box
box1cnt - counter to determine the number of consecutive points having the same
amplitude in box 1
box1Aold - Previous amplitude value in box 1
box2cnt - counter to determine the number of consecutive points having the same
amplitude in box 1
box2Aold - Previous amplitude value in box 2
box3cnt - counter to determine the number of consecutive points having the same
amplitude in box 1
box3Aold - Previous amplitude value in box 3

```

These parameters correspond to the data spaces, reference values, and counts necessary for providing the RUL estimate.

The second UPS being monitored has an increasing load placed on its output. The SoH output is shown in Figure 9. This screen shot demonstrates the UPS state of health adjusting directly with the AC output load. Recall that while this is a direct relationship, the plot depicts remaining health.

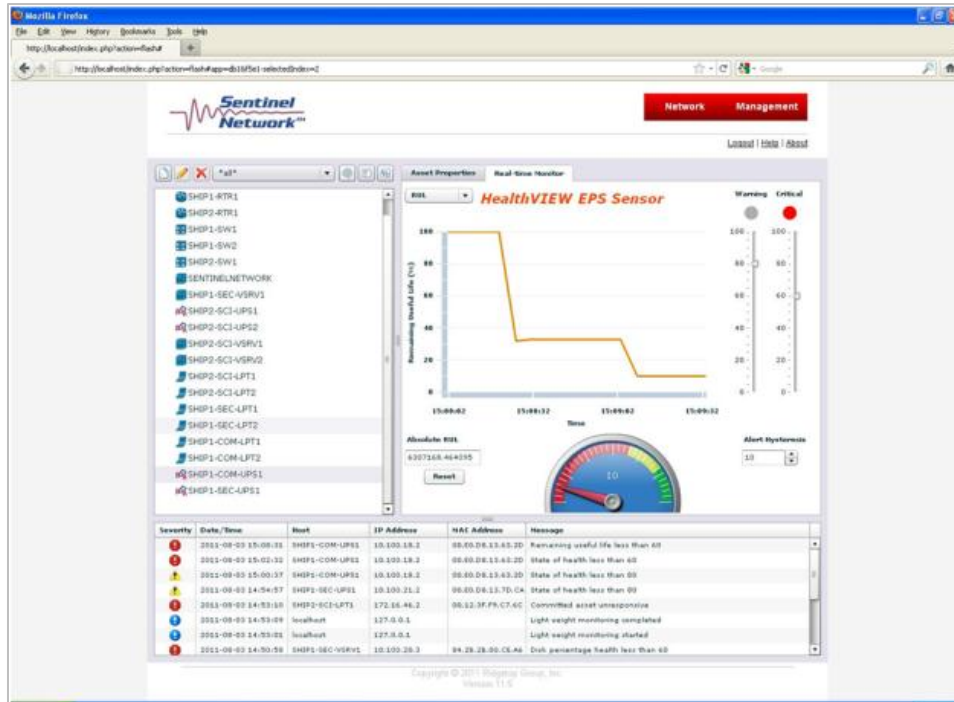


Figure 9: RUL output for a UPS with an increasing load supplied

Background monitoring of the UPS devices consists of configuring the setting for the interval (1 minute), warning threshold percent (80%), critical threshold percent (60%), hysteresis percent (10%), and the IP addresses of the devices being monitored. In this case the IP addresses corresponding to the two UPS devices monitored during real-time monitoring were selected. Background data collection consists of polling the device every minute through SNMP, retrieving data from the UPS, and converting that data into an SOH and RUL estimate. This information is stored in a database and can be displayed on-screen using filters to display by certain device, information type, and value.

Conclusion and Future Developments

For further development on the product roadmap, Ridgetop will continue to develop and refine sensors that can be integrated with existing UPS firmware; see Figure 10. This sensor converts data collected from the UPS into an SOH reading. Sentinel Network will convert the SOH into an RUL value and provide troubleshooting support when failure has been predicted.

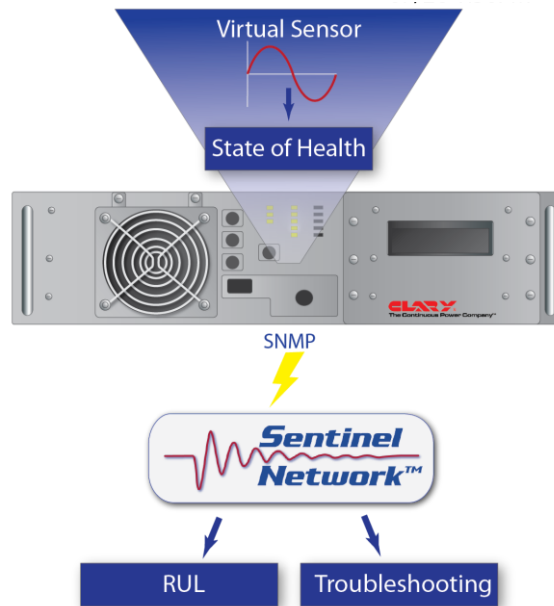


Figure 10: Sensor integration

With an accurate RUL, condition-based maintenance (CBM) can be supported, reducing network downtime. In this manner an overlooked niche in the market can be supported by incorporating prognostics into IT monitoring and management software tools. Reasoners can provide a competitive advantage to any IT partner, through the capability of generating reports of actionable items to solve network failures before they occur.