Network Working Group B. Claise
Internet-Draft J. Parello
Intended Status: Informational B. Schoening
Expires: April 20, 2011 Cisco Systems, Inc.
October 20, 2010

Power Management Architecture
draft-claise-power-management-arch-02

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on September, 2010.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

This document defines the power management architecture.

Table of Contents

TO DO

- Question for the Working Group: Should the WG consider IPFIX in this architecture?

- Question for the Working Group: How to specify the notion of child capabilities, i.e. the capabilities that the Power Monitor Parents have with Power Monitor Children. For Example:
   1. Monitoring (only reporting)
   2. Configuration power state
   3. Configuration: power

  Example: on a PC, we can set power level without knowing the power. A solution must be specified in this draft.

- Question for the Working Group: Should transition states be tracked when setting a level. Example: The configured level is set to Off from High. The Actual level will take time to update as the device powers down. Should there be transitions shown or will the two variables suffice to track the device state.

- Question for working group: Should implementation scenarios be incorporated in the architecture draft

- We should have a similar section, for all the drafts, which includes an overview of all EMAN documents.

1. Introduction

Network management is typically divided into the five main network management areas defined in the ISO Telecommunications Management Network model: Fault, Configuration, Accounting, Performance, and Security Management. Absent from this model is any consideration of energy management, which is now becoming a critical area of concern worldwide.

This document defines an architecture for power management for devices within or connected to communication networks. This architecture includes monitoring for power state and energy consumption of networked elements, covering the requirements specified in [POWER-MON-REQ]. It also goes a step further in defining some elements of configuration.

Energy management is applicable to devices in communication networks. Target devices for this specification include (but are not limited to): routers, switches, Power over Ethernet (PoE) endpoints, protocol gateways for building management systems, intelligent meters, home energy gateway, hosts and servers, sensor proxies, etc.

Where applicable, device monitoring extends to the individual components of the device and to any attached dependent devices. For example: A device can contain components that are independent from a power-state point of view, such as line cards, processor cards, hard drives. A device can also have dependent attached devices, such as a switch with PoE endpoints or a power distribution unit with attached endpoints.

2. Use Cases & Requirements

Requirements for power and energy monitoring for networking devices are specified in [POWER-MON-REQ]. The requirements in [POWER-MON-REQ] cover devices typically found in communications networks, such as switches, routers, and various connected endpoints. For a power monitoring architecture to be useful, it should also apply to facility meters, power distribution units, gateway proxies for commercial building control, home automation devices, and devices that interface with the utility and/or smart grid. Accordingly, this architecture, the scope is broader than that specified in [POWER-MON-REQ]. Several scenarios that cover these broader use cases are presented later in Section 11. - Implementation Scenarios.

3. Terminology

This section contains definitions of important terms used throughout this specification.

IPFIX-specific terminology used in this document is defined in section 2 of [RFC5101]. For example: Flow Record, Collector , etc... As in [RFC5101], these IPFIX-specific terms have the first letter of a word capitalized.

Power Monitor

A Power Monitor is a component within a system of components that provides power, draws power, or reports energy consumption on behalf of another Power Monitor. It can be independently

managed from a power-monitoring and power-state configuration point of view. Examples of Power Monitors are: a router line card, a motherboard with a CPU, an IP phone connected with a switch, etc.

Power Monitor Parent

A Power Monitor Parent is a Power Monitor that is the root of one or more subtending Power Monitors, called Power Monitor Children. The Power Monitor Parent is able to collect data about or report on the power state and energy consumption of its Power Monitor Children.

For example: A Power-over-Ethernet (PoE) device (such as an IP phone or an access point) is attached to a switch port. The switch is the source of power for the attached device, so the Power Monitor Parent is the switch, and the Power Monitor Child is the device attached to the switch.

The Power Monitor Parent may report data or implement actions on behalf of the Power Monitor Child. These capabilities must be enumerated by the Power Monitor Parent.

The communication between the parent and child for monitoring or collection of power data is left to the device manufacturer. For example: A parent switch may use LLDP to communicate with a connected child, and a parent lighting controller may use BACNET to communicate with child lighting devices.

Power Monitor Child

A Power Monitor Child is a Power Monitor associated with a Power Monitor Parent, and which reports its power usage and power state to its Power Monitor Parent. The Power Monitor Child may or may not draw power from its Power Monitor Parent. .

Power Monitor Meter Domain

A Power Monitor Meter Domain is a name or name space that logically groups Power Monitors into a zone of manageable power usage. Typically, this zone will have as members all Power Monitors that are powered from the same electrical panel or panels for which there is a meter or sub meter. For example: All Power Monitors receiving power from the same distribution panel of a building, or all Power Monitors in a building for

which there is one main meter, would comprise a Power Monitor Meter Doman. From the standpoint of power-use monitoring, it is useful to report the total power usage as the sum of power consumed by all the Power Monitors within a Power Monitor Meter Domain and then correlate that value with the metered usage.

Power Level

A Power Level is a uniform way to classify power settings on a Power Monitor (e.g., shut, hibernate, sleep, high). Power Levels can be viewed as an interface for the underlying device-implemented power settings.

Manufacturer Power Level

A Manufacturer Power Level is a device-specific way to classify power settings implemented on a Power Monitor. For cases where the implemented power settings cannot be directly mapped to Power Levels, we can use the Manufacturer Power Levels to enumerate and show the relationship between the implemented power settings and the Power Level interface.

4. Energy Management Reference Model

```
                  +----------------+
                  |      NMS       |            -
                  +-----+---+------+            |
                        |   |                   |
                        |   |                   |  S
              +---------+   +-------+           |  N
              |                     |           |  M
              |                     |           |  P
      +----------------+      +------+-------+     |
      | Power Monitor |       | Power Monitor |    |
      | Parent 1      | ...   | Parent N      |    -
      +----------------+      +----------------+
              |||
    (protocol |||
    out of    |||      +-------------+---------+
    the scope)|||------| Power Monitor Child 1 |
              ||       +-----------------------+
              ||
              ||       +-------------+---------+
              ||-------| Power Monitor Child 2 |
```

```
              |        +-----------------------+
              |
              |
              |--------           ...
              |
              |
              |        +-------------+---------+
              |--------| Power Monitor Child M |
                       +-----------------------+
```

Figure 1: Energy Management Pull Reference Model

In this architecture a Network Management Station (NMS) will poll MIB variables on a Power Monitors via SNMP. The Power Monitor returns information for itself and for any Power Monitor Children if applicable. The information returned will contain business context, energy usage, power quality and other information as described further.

The protocol between the Power Monitor Parent and Power Monitor Children is out of scope of this document. The Power Monitor Parent may speak to a Power Monitor Child using a manufacturer selected protocol. This protocol may or may not based on IP. In this way, a Power Monitor Parent acts as a PROXY for protocol translation between the Power Monitor Parent and Child. The Power Monitor Parent also acts as an aggregation point for other subtended Power Monitor Children.

```
                  +---------------+
                  | NMS/Collector |              ^  S
                  +-----+---+-----+              |  N
                        |  |                     |  M     I
                        |  |                     |  P     P
              +---------+  +-------+             |     &  F
              |                     |            |  N     I
              |                     |            |  O     X
      +---------------+      +------+-------+    |  T
      | Power Monitor |      | Power Monitor |   |  .
      | Parent 1      | ...  | Parent N      |   -
      +---------------+      +---------------+
              |||
    (protocol |||
    out of    |||      +-------------+---------+
    the scope)|||------| Power Monitor Child 1 |
              ||       +-----------------------+
              ||
```

```
              ||       +-------------+---------+
              ||-------| Power Monitor Child 2 |
              |        +-----------------------+
              |
              |
              |--------           ...
              |
              |
              |        +-------------+---------+
              |--------| Power Monitor Child M |
                       +-----------------------+
```

Figure 2: Energy Management PUSH Reference Model

The Power Monitor Parents may send SNMP notifications regarding their own state or the state of their Power Monitor Children. The Power Monitor Children do not send SNMP notifications on their own.

As discussed in [POWER-MON-REQ], the Power Monitor Parents may export IPFIX Flow Records [RFC5101] to a Collector. The IPFIX protocol is well suited for regular time series export of similar information, such as the energy consumed by the Power Monitor Children.

EDITOR'S NOTE: at this point in time, there is no draft specifying the IPFIX Flow Records.

5. Architecture High Level Concepts and Scope

The scope of this architecture is to enable networking and network-attached devices to be managed with respect to their energy consumption or production. The goal is to make devices energy-aware.

The architecture describes how to make a device aware of its consumption or production of energy expressed as usage in watts. This does not include:

- Manufacturing costs in currency or environmental units
- Embedded carbon or environmental equivalences of the device itself
- Cost in currency or environmental impact to dismantle or recycle the device
- Relationship to an electrical or smart grid
- Supply chain analysis

- Conversion of the usage or production of energy to units expressed from the source of that energy (such as the greenhouse gas emissions associated with 1000kW from a diesel source).

The remainder of this section describes the basic concepts of the architecture. Each concept is examined in detail in subsequent sections.

Examples are provided in a later section to show how these concepts can be implemented.

The basic concepts are:

The Power Monitor will have basic naming and informational descriptors to identify it in the network.

A Power Monitor can be part of a Power Monitor Meter Domain. A Power Monitor Meter Domain is a manageable set of devices that has a meter or sub-meter attached and typically corresponds to a power distribution point or panel.

A Power Monitor can be a parent (Power Monitor Parent) or child (Power Monitor Child) of another Power Monitor. This allows for Power Monitor Parent to aggregate power reporting and control of power information.

Each Power Monitor can have information to allow it to be described in the context of the business or ultimate use. This is in addition to its networked information. This allows for tagging, grouping, and differentiation between Power Monitors for NMS.

For control and universal monitoring, each Power Monitor implements or declares a set of known Power Levels. The Power Levels are mapped to Manufacturer Power Levels that indicate the specific power settings for the device implementing the Power Monitor.

When the Power Level is set, a Power Monitor may be busy at the request time. The Power Monitor will set the desired level and then update the actual Power Level when the priority task is finished. This mechanism implies two different Power Level variables: actual versus desired.

EDITOR'S NOTE: The transition state will have to be specified.

Each Power Monitor will have usage information that describes the power information along with how that usage was obtained or derived.

Optionally, a Power Monitor can further describe the power information with power quality information reflecting the electrical characteristics of the measurement.

Optionally, a Power Monitor can provide power usage over time to describe energy consumption

If a Power Monitor has one or more batteries, it can provide optional battery information as well.

5.1. Power Monitor Information

Every Power Monitor should have a unique printable name, and must have a unique Power Monitor index.

Possible naming conventions are: textual DNS name, MAC-address of the device, interface ifName, or a text string uniquely identifying the Power Monitor. As an example, in the case of IP phones, the Power Monitor name can be the device DNS name.

5.2. Power Monitor Meter Domain

Each Power Monitor must be a member of a Power Monitor Meter Domain. The Power Monitor Meter Domain should map 1-1 with a metered or sub-metered portion of the site. The Power Monitor Meter Domain must be configured on the Power Monitor Parent. The Power Monitor Children may inherit their domain values from the Power Monitor Parent or the Power Monitor Meter Domain may be configured directly in a Power Monitor Child.

5.3. Power Monitor Parent and Child

A Power Monitor Child reports its power usage to its Power Monitor Parent. A Power Monitor Child has one and only one Power Monitor Parent. If a Power Monitor had two parents there would be a risk of double-reporting the power usage in the Power Monitor Meter Domain. Therefore, a Power Monitor cannot be both a Power Monitor Parent and a Power Monitor Child at the same time.

A Power Monitor Child can be fully dependent on the Power Monitor Parent for its power or independent from the parent (such as a PC connected to a switch). In the dependently powered case, the Power Monitor Parent provides power for the Power Monitor Child (as in the case of Power Over Ethernet devices). In the independently powered case, the Power Monitor Child draws power from another source (typically a wall outlet). Since the Power Monitor Parent is not the source of power supply, the power usage cannot be measured at the Power Monitor Parent. However, an independent Power Monitor Child reports Power Monitor information to the Power Monitor Parent. The Power Monitor Child may listen to the power control settings from a Power Monitor Parent and could react to the control messages. However, note that the communication between the Power Monitor Parent and Power Monitor Child is out of scope for this document.

A mechanism, outside of the scope of this document, should be in place to verify the connectivity between the Power Monitor Parent and its Power Monitor Children. If a Power Monitor Child is unavailable, the Power Monitor Parent must follow some rules to determine how long it should wait before removing the Power Monitor Child entry, along with all associated statistics, from its database. In some situations, such as a connected building in which the Power Monitor Children are somewhat static, this removal-delay period may be long, and persistence across a Power Monitor Parent reload may make sense. However, in a networking environment, where endpoints can come and go, there is not much sense in configuring a long removal timer. In all cases, the removal timer or persistence must be clearly specified.

Further examples of Power Monitor Parent and Child implementations are provided in the Implementation Scenarios section 11.

5.4. Power Monitor Context

Monitored power data will ultimately be collected by and reported from an NMS. In order to aid in reporting and in differentiation between Power Monitors, each Power Monitor will contain information establishing its business or site context. A Power Monitor can provide an importance value in the range of 1 to 100 to help differentiate a device's use or relative value to the site. The importance range is from 1 (least important) to 100 (most important). The default importance value is 1.

For example: A typical office environment has several types of phones, which can be rated according to their business impact. A public desk phone has a lower importance (for example, 10) than a business-critical emergency phone (for example, 100). As another example: A company can consider that a PC and a phone for a customer-service engineer is more important than a PC and a phone for lobby use.

Although network managers must establish their own ranking, the following is a broad recommendation:

. 90 to 100 Emergency response
. 80 to 90 Executive or business-critical
. 70 to 79 General or Average
. 60 to 69 Staff or support
. 40 to 59 Public or guest
. 1 to 39 Decorative or hospitality

A Power Monitor can provide a set of keywords. These keywords are a list of tags that can be used for grouping and summary reporting within or between Power Monitor Meter Domains. All alphanumeric characters and symbols, such as #, (, $, !, and &, are allowed. Potential examples are: IT, lobby, HumanResources, Accounting, StoreRoom, CustomerSpace, router, phone, floor2, or SoftwareLab. There is no default value for a keyword.

Multiple keywords can be assigned to a device. In such cases, the keywords are separated by commas and no spaces between keywords are allowed. For example, "HR,Bldg1,Private".

Additionally, a Power Monitor can provide a "role description" string that indicates the purpose the Power Monitor serves in the network or for the site/business. This could be a string describing the context the device fulfills in deployment. For example, a lighting fixture in a kitchen area could have a role of "Hospitality Lighting" to provide context for the use of the device.

5.5. Power Monitor Levels

Power Levels represent universal states of power management of a Power Monitor. Each Power Level corresponds to a global, system, and performance state in the ACPI model [ACPI].

Level	ACPI Global/System State	Power Level Name

Non-operational states:

1	G3, S5	Mech Off
2	G2, S5	Soft Off
3	G1, S4	Hibernate
4	G1, S3	Sleep
5	G1, S2	Standby
6	G1, S1	Ready

Operational states:

7	G0, S0, P5	LowMinus
8	G0, S0, P4	Low
9	G0, S0, P3	MediumMinus
10	G0, S0, P2	Medium
11	G0, S0, P1	HighMinus
12	G0, S0, P0	High

Figure 3: ACPI / Power Level Mapping

For example, a Power Monitor with a Power Level of 9 would indicate an operational state with MediumMinus Power Level.

The Power Levels can be considered as guidelines in order to promote interoperability across device types. Realistically, each specific feature requiring Power Levels will require a complete recommendation of its own. For example, designing IP phones with consistent Power Levels across vendors requires a specification for IP phone design, along with the Power Levels mapping.

Manufacturer Power Levels are required in some situations, such as when no mappings with the existing Power Levels are possible, or when more than the twelve specified Power Levels are required.

A first example would be an imaginary device type, with only five levels: "none", "short", "tall", "grande", and "venti".

Manufacturer Power Level	Respective Name
0	none
1	short
2	tall
3	grande
4	venti

Figure 4: Mapping Example 1

In the unlikely event that there is no possible mapping between these Manufacturer Power Levels and the proposed Power Monitor Power Levels, the Power Level will remain 0 throughout the MIB module, as displayed below.

Power Level / Name	Manufacturer Power Level / Name
0 / unknown	0 / none
0 / unknown	1 / short
0 / unknown	2 / tall
0 / unknown	3 / grande
0 / unknown	4 / venti

Figure 5: Mapping Example 2

If a mapping between the Manufacturer Power Levels and the Power Monitor Power Levels is achievable, both series of levels must exist in the MIB module in the Power Monitor Parent, allowing the NMS to understand the mapping between them by correlating the Power Level with the Manufacturer Power Levels.

Power Level / Name	Manufacturer Power Level / Name
1 / Mech Off	0 / none
2 / Soft Off	0 / none
3 / Hibernate	0 / none
4 / Sleep, Save-to-RAM	0 / none
5 / Standby	0 / none
6 / Ready	1 / short
7 / LowMinus	1 / short
8 / Low	1 / short
9 / MediumMinus	2 / tall
10 / Medium	2 / tall
11 / HighMinus	3 / grande
12 / High	4 / venti

Figure 6: Mapping Example 3

How the Power Monitor Levels are then mapped is an implementation choice. However, it is recommended that the Manufacturer Power Levels map to the lowest applicable Power Levels, so that setting all Power Monitors to a Power Level would be conservative in terms of disabled functionality on the Power Monitor.

A second example would be a device type, such as a dimmer or a motor, with a high number of operational levels. For the sake of the example, 100 operational states are assumed.

```
    Power Level / Name       Manufacturer Power Level / Name
        1 / Mech Off                  0 / off
        2 / Soft Off                  0 / off
        3 / Hibernate                 0 / off
        4 / Sleep, Save-to-RAM        0 / off
        5 / Standby                   1 / off
        6 / Ready                     2 / off
        7 / LowMinus                  11 / 1%
        7 / LowMinus                  12 / 2%
        7 / LowMinus                  13 / 3%
        .                             .
        .                             .
        .                             .
        8 / Low                       15 / 15%
        8 / Low                       16 / 16%
        8 / Low                       17 / 17%
        .                             .
        .                             .
        .                             .
        9 / MediumMinus               30 / 30%
        9 / MediumMinus               31 / 31%
        9 / MediumMinus               32 / 32%
        .                             .
        .                             .
        .                             .
        10 / Medium                   45 / 45%
        10 / Medium                   46 / 46%
        10 / Medium                   47 / 47%
        .                             .
        .                             .
        .                             .
        etc...
```

Figure 7: Mapping Example 4

As specified in section 6, this architecture allows the configuration of the Power Level, while configuring the Manufacturer Power Level from the MIB directly is not possible.

5.6. Power Monitor Usage Measurement

The usage or production or power must be qualified as more than a value alone. A measurement should be qualified with the units, magnitude, direction of power flow, and by what means the measurement was made (ex: Root Mean Square versus Nameplate) .

In addition, the Power Monitor should describe how it intends to measure usage as one of consumer, producer or meter of usage. Given the intent any readings can be correctly summarized or analyzed by an NMS. For example metered usage reported by a meter and consumption usage reported by a device connected to that meter may naturally measure the same usage. With the two measurements identified by intent a proper summarization can be made by an NMS.

The power usage measurement should conform to the IEC 61850 definition of unit multiplier for the SI (System International) units of measure. The power usage measurement is considered an instantaneous usage value and does not include the usage over time.

Measured values are represented in SI units obtained by BaseValue * 10 raised to the power of the scale. For example, if current power usage of a Power Monitor is 3, it could be 3 W, 3 mW, 3 KW, or 3 MW, depending on the value of the scaling factor

In addition to knowing the usage and magnitude, it is useful to know how a Power Monitor usage measurement was obtained:

- Whether the measurements were made at the device itself or from a remote source.
- Description of the method that was used to measure the power and whether this method can distinguish actual or estimated values.

An NMS can use this information to account for the accuracy and nature of the reading between different implementations.

In addition to the power usage, the nameplate power rating of a Power Monitor is typically specified by the vendor as the capacity required to power the device. Often this label is a conservative number and is the worst-case power draw. While the actual utilization of an entity can be lower, the nameplate power is important for provisioning, capacity planning and billing.

5.7. Optional Power Usage Quality

Given a power measurement of a Power Monitor, it may in certain circumstances be desirable to know the power quality associated with that measurement. The information model must adhere to the IEC 61850 7-2 standard for describing AC measurements. In some

Power Monitor Domains, the power quality may not be needed, available, or relevant to the Power Monitor.

5.8. Optional Energy Measurement

In addition to reporting the Power Level, an approach to characterizing the energy demand is required. It is well known in commercial electrical utility rates that demand charges can be on par with actual power charges, so it is useful to characterize the demand. The demand can be described as the average energy of an Power Monitor over a time window called a demand interval (typically 15 minutes). The highest peak energy demand measured over a time horizon, such as 1 month or 1 year, is often the basis for usage charges. A single window of time of high usage can penalize the consumer with higher energy consumption charges. However, it is relevant to measure the demand only when there are actual power measurements from a Power Monitor, and not when the power measurement is assumed or predicted.

Several efficiency metrics can be derived and tracked with the demand usage data. For example:

- Per-packet power costs for a networking device (router or switch) can be calculated by an NMS. The packet count can be determined from the traffic usage in the ifTable [RFC2863], from the forwarding plane figure, or from the platform specifications.

- Watt-hour power can be combined with utility energy sources to estimate carbon footprint and other emission statistics.

5.9. Optional Battery Information

Some Power Monitors may be running on batteries. Therefore information such as the battery status (charging or discharging), remaining capacity, and so on, must be available.

6. Power Monitor Children Discovery

There are multiple ways that the Power Monitor Parent can discover its Power Monitor Children, if they are not present on the same physical network element:

. In case of PoE, the Power Monitor Parent automatically discovers a Power Monitor Child when the Child requests power.
. The Power Monitor Parent and Children may run the Link Layer Discovery Protocol [LLDP], or any other discovery protocol, such as Cisco Discovery Protocol (CDP). The Power Monitor Parent might even support the LLDP-MED MIB [LLDP-MED-MIB], which returns extra information on the Power Monitor Children.
. The Power Monitor Parent may reside on a network connected facilities gateway. A typical example is a converged building gateway, monitoring several other devices in the building, and serving as a proxy between SNMP and a protocol such as BACNET.

When a Power Monitor Child supports only its own Manufacturer Power Levels, the Power Monitor Parent will have to discover those Manufacturer Power Levels. Note that the communication specifications between the Power Monitor Parent and Children is out of the scope of this document. This includes the Manufacturer Power Levels discovery, which is protocol-specific.

7. Configuration

This power management architecture allows the configuration of the following key parameters:

. Power Monitor name: A unique printable name for the Power Monitor.
. Power Monitor Role: An administratively assigned name to indicate the purpose a Power Monitor serves in the network.
. Power Monitor Importance: A ranking of how important the Power Monitor is, on a scale of 1 to 100, compared with other Power Monitors in the same Power Monitor Meter Domain.
. Power Monitor Keywords: A list of keywords that can be used to group Power Monitors for reporting or searching.
. Power Monitor Domain: Specifies the name of a Power Monitor Meter Domain for the Power Monitor.
. The Power Monitor Level: Specifies the current Power Level (0..12) for the Power Monitor.
. The energy demand parameters: For example, which interval length to report the energy on, the number of intervals to keep, etc.

When a Power Monitor requires a mapping with the Manufacturer Power Level, the Power Monitor configuration is done via the

Power Level settings, and not directly via the Manufacturer Power Levels, which are read-only. Taking into account Figure 8, where the LowMinus Power Level corresponds to three different Manufacturer Power Levels (11 for 1%, 12 for 2%, and 13 for 3%), the implication is that this architecture will not set the Manufacturer Power Level to one percent granularity without communicating over or configuring the proprietary protocol for this Power Monitor.

This architecture uses a Power Level MIB object to set up the Power Level for a specific Power Monitor. However, the Power Monitor might be busy executing an important task that requires the current Power Level for some more time. For example, a PC might have to finish a backup first, or an IP phone might be busy with a current phone call. Therefore a second MIB object contains the actual Power Level. A difference in values between the two objects indicates that the Power Monitor is currently in Power Level transition.

Interactions with established open protocols, such as Wake-up-on-Lan (WoL) and DASH [DASH], may require configuration in the Power Monitor as well, facilitating the communication between Power Monitor Parent and remote Power Monitor Children.

Note that the communication specifications between the Power Monitor Parent and Children is out of the scope of this document. This includes communication of power settings and configuration information, such as the Power Monitor Domain.

8. Fault Management

[POWER-MON-REQ] specifies some requirements about power states such as "the current state - the time of the last change", "the total time spent in each state", "the number of transitions to each state", etc. Such requirements are fulfilled via the pmPowerLevelChange NOTIFICATION-TYPE [POWER-MON-MIB]. This SNMP notification is generated when the value(s) of Power Level has changed for the Power Monitor.

9. IPFIX

A push-based mechanism, such as IPFIX [RFC5101], might be required to export high-volume time series of energy consumption values, as mentioned in [POWER-MON-REQ].

EDITOR'S NOTE: the Working Group should decide how much of IPFIX should be described in this document

10. Relationship with Other Standards Development Organizations

10.1. Information Modeling

This power management architecture should, as much as possible, reuse existing standards efforts, especially with respect to information modeling and data modeling [RFC3444].

The data model for power, energy related objects is based on IEC 61850.

Specific examples include:

- The scaling factor, which represents Power Monitor usage magnitude, conforms to the IEC 61850 definition of unit multiplier for the SI (System International) units of measure.

- The power accuracy model is based on the ANSI and IEC Standards, which require that we use an accuracy class for power measurement. ANSI and IEC define the following accuracy classes for power measurement:

  - IEC 62053-22 60044-1 class 0.1, 0.2, 0.5, 1 3.

  - ANSI C12.20 class 0.2, 0.5

- The powerQualityMIB MIB module adheres closely to the IEC 61850 7-2 standard for describing AC measurements.

10.2. Power Levels

There are twelve Power Monitor Levels. They are subdivided into six operational states, and six non-operational states. The lowest non-operational state is 1 and the highest is six. Each non-operational state corresponds to an ACPI level [ACPI].

11. Implementation Scenarios

The scope of power and energy monitoring consists of devices that consume power within and that are connected to a communications network. These devices include:

- Network devices and sub-components: Devices such as routers and switches and their sub-components.

- Network attached endpoints: Devices that use the communications network, such as endpoints, PCs, and facility gateways that proxy energy monitor and control for commercial buildings or home automation.

- Network attached meters or supplies: Devices that can monitor the electrical supply, such as smart meters or Universal Power Supplies (UPS) that meter and provide availability.
-

This section provides illustrative examples that model different scenarios for implementation of the Power Monitor, including Power Monitor Parent and Power Monitor Child relationships.

Each of the scenarios below is explained in more detail in the Power Monitor MIB document [POWER-MON-MIB], with a mapping to the MIB Objects.

Scenario 1: Switch with PoE endpoints

Consider a PoE IP phone connected to a switch. The IP phone draws power from the PoE switch.

Scenario 2: Switch with PoE endpoints with further connected device(s)

Consider the same example as in Scenario 1, but with a PC daisy-chained from the IP phone for LAN connectivity. The phone draws power from the PoE port of the switch, while the PC draws power from the wall outlet.

Scenario 3: A switch with Wireless Access Points

Consider a WAP (Wireless Access Point) connected to the PoE port of a switch. There are several PCs connected to the Wireless Access Point over Wireless protocols. All PCs draw power from the wall outlets.

The switch port is the Power Monitor Parent for the Wireless Access Point (WAP) and all the PCs. But there is a distinction among the Power Monitor Children, as the WAP draws power from the PoE port of the switch and the PCs draw power from the wall outlet.

Scenario 4: Network connected facilities gateway

At the top of the network hierarchy of a building network is a gateway device that can perform protocol conversion between many facility management devices, such as BACNET, MODBUS, DALI, LON, etc. There are power meters associated with power-consuming entities (Heating Ventilation & Air Conditioning - HVAC, lighting, electrical, fire control, elevators, etc). The proposed MIB can be implemented on the gateway device. The gateway can be considered as the Power Monitor Parent, while the power meters associated with the energy consuming entities can be considered as its Power Monitor Children.

Scenario 5: Data center network

A typical data center network consists of a hierarchy of switches. At the bottom of the hierarchy there are servers mounted on a rack, and these are connected to the top-of-the-rack switches. The top switches are connected to aggregation switches that are in turn connected to core switches. As an example, Server 1 and Server 2 are connected to different switch ports of the top switch.

The proposed MIB can be implemented on the switches. The switch can be considered as the Power Monitor Parent. The servers can be considered as the Power Monitor Children.

Scenario 6: Building gateway device

Similar scenario as the scenario 4.

Scenario 7: Power consumption of UPS

Data centers and commercial buildings can have Uninterruptible Power Supplies (UPS) connected to the network. The Power Monitor can be used to model a UPS as a Power Monitor Parent with the connected devices as Power Monitor Children.

Scenario 8: Power consumption of battery-based devices

A PC is a typical example of a battery-based device.

12. Security Considerations

Regarding the data attributes specified here, some or all may be considered sensitive or vulnerable in some network environments. Reading or writing these attributes without proper protection such as encryption or access authorization may have negative effects on the network capabilities.

12.1. Security Considerations for SNMP

Readable objects in a MIB modules (i.e., objects with a MAX-ACCESS other than not-accessible) may be considered sensitive or vulnerable in some network environments. It is thus important to control GET and/or NOTIFY access to these objects and possibly to encrypt the values of these objects when sending them over the network via SNMP.

The support for SET operations in a non-secure environment without proper protection can have a negative effect on network operations. For example:

- Unauthorized changes to the Power Domain or business context of a Power Monitor may result in misreporting or interruption of power.
- Unauthorized changes to a power level may disrupt the power settings of the different Power Monitors, and therefore the level of functionality of the respective Power Monitors.
- Unauthorized changes to the demand history may disrupt proper accounting of energy usage.

With respect to data transport SNMP versions prior to SNMPv3 did not include adequate security. Even if the network itself is secure (for example, by using IPsec), there is still no secure control over who on the secure network is allowed to access and GET/SET (read/change/create/delete) the objects in these MIB modules.

It is recommended that implementers consider the security features as provided by the SNMPv3 framework (see [RFC3410], section 8), including full support for the SNMPv3 cryptographic mechanisms (for authentication and privacy).

Further, deployment of SNMP versions prior to SNMPv3 is not recommended. Instead, it is recommended to deploy SNMPv3 and to enable cryptographic security. It is then a customer/operator responsibility to ensure that the SNMP entity giving access to an instance of these MIB modules is properly configured to give access to the objects only to those principals (users) that have legitimate rights to GET or SET (change/create/delete) them.

12.2. Security Considerations for IPFIX

EDITOR'S NOTE: to be completed if IPFIX is discussed in this document

13. IANA Considerations

This document has no actions for IANA.

14. Acknowledgments

The authors would like to Michael Brown for improving the text dramatically.

15. References

Normative References

[RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet Standard Management Framework ", RFC 3410, December 2002.

[RFC5101] B. Claise, Ed., Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information, RFC 5101, January 2008.

[POWER-MON-REQ] Quittek, J., Winter, R., Dietz, T., Claise, B., and M. Chandramouli, "Requirements for Power Monitoring", draft-quittek-power-monitoring-requirements-01 (work in progress), July 2010.

[POWER-MON-MIB] Claise, B., Chandramouli, M., Parello, J., and Schoening, B., "Power and Energy Monitoring MIB", draft-claise-energy-monitoring-mib-06, (work in progress), October 2010.

Informative References

[RFC2863] McCloghrie, K., Kastenholz, F., "The Interfaces Group MIB", RFC 2863, June 2000.

[RFC3444] Pras, A., Schoenwaelder, J. "On the Differences between Information Models and Data Models", RFC 3444, January 2003.

[ACPI] "Advanced Configuration and Power Interface Specification", http://www.acpi.info/spec30b.htm

[LLDP] IEEE Std 802.1AB, "Station and Media Control Connectivity Discovery", 2005.

[LLDP-MED-MIB] ANSI/TIA-1057, "The LLDP Management Information Base extension module for TIA-TR41.4 media endpoint discovery information", July 2005.

[DASH] "Desktop and mobile Architecture for System Hardware", http://www.dmtf.org/standards/mgmt/dash/

Authors' Addresses

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Diegem 1813
BE

Phone: +32 2 704 5622
Email: bclaise@cisco.com

John Parello
Cisco Systems, Inc.
3550 Cisco Way
San Jose, California 95134
US

Phone: +1 408 525 2339
Email: jparello@cisco.com

Brad Schoening
Cisco Systems, Inc.
3550 Cisco Way
San Jose, California 95134
US

Phone: +1 408 525 2339
Email: braschoe@cisco.com