

Network Working Group
Internet-Draft
Intended Status: Standards Track
Expires: January 8 2012

M. Chandramouli
Cisco Systems, Inc.
B. Schoening
Independent Consultant
J. Quittek
T. Dietz
NEC Europe Ltd.
B. Claise
Cisco Systems, Inc.
July 8, 2011

Power and Energy Monitoring MIB
draft-claise-energy-monitoring-mib-09

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/lid-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 January 2012.

<Claise, et. Al>

Expires January 8, 2012

[Page 1]

Copyright (c) 2011 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 a subset of the Management Information Base (MIB) for power and energy monitoring of devices.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction.....	3
2. The Internet-Standard Management Framework.....	4
3. Use Cases.....	4
4. Terminology.....	5
5. Architecture Concepts Applied to the MIB Module.....	5
5.1. Power Monitor Information.....	11
5.2. Power State.....	12
5.2.1. Power State Set.....	13
5.2.2. IEEE1621 Power State Set.....	13
5.2.3. DMTF Power State Set.....	13
5.2.4. EMAN Power State Set.....	14
5.3. Power Monitor Usage Information.....	17
5.4. Optional Power Usage Quality.....	18
5.5. Optional Energy Measurement.....	19

Internet-Draft	<Energy Monitoring MIB>	July 2011
5.6. Fault Management.....		22
6. Discovery.....		23
6.1. ENERGY-AWARE-MIB Module Implemented.....		23
6.2. ENERGY-AWARE-MIB Module Not Implemented, ENTITY-MIB Implemented.....		24
6.3. ENERGY-AWARE-MIB Module and ENTITY-MIB Not Implemented..		24
7. Link with the other IETF MIBs.....		24
7.1. Link with the ENTITY MIB and the ENTITY-SENSOR MIB..		24
7.2. Link with the ENTITY-STATE MIB.....		26
7.3. Link with the POWER-OVER-ETHERNET MIB.....		26
7.4. Link with the UPS MIB.....		27
7.5. Link with the LLDP and LLDP-MED MIBs.....		28
8. Implementation Scenarios.....		29
9. Structure of the MIB.....		31
10. MIB Definitions.....		31
11. Security Considerations.....		66
12. IANA Considerations.....		67
12.1. IANA Considerations for the MIB Modules.....		67
12.2. IANA Registration of new Power State Set.....		67
12.2.1. IANA Registration of the IEEE1621 Power State Set..		68
12.2.2. IANA Registration of the DMTF Power State Set.....		68
12.2.3. IANA Registration of the EMAN Power State Set.....		69
12. Contributors.....		69
13. Acknowledgment.....		69
14. Open Issues.....		69
15. References.....		71
15.2. Normative References.....		71
15.3. Informative References.....		71

1. Introduction

This document defines a subset of the Management Information Base (MIB) for use in energy management of devices within or connected to communication networks. The MIB modules in this document are designed to provide a model for energy management, which includes monitoring for power state and energy consumption of networked elements. This MIB takes into account the Power Management Architecture [EMAN-FRAMEWORK], which in turn, is based on the Power Monitoring Requirements [EMAN-REQ].

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

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.

Devices and their sub-components may be characterized by the power-related attributes of a physical entity present in the ENTITY MIB, even though the ENTITY MIB compliance is not a requirement due to the variety and broad base of devices concerned with energy management.

2. The Internet-Standard Management Framework

For a detailed overview of the documents that describe the current Internet-Standard Management Framework, please refer to section 7 of RFC 3410 [RFC3410].

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB are defined using the mechanisms defined in the Structure of Management Information (SMI). This memo specifies MIB modules that are compliant to SMIv2, which is described in STD 58, RFC 2578 [RFC2578], STD 58, RFC 2579 [RFC2579] and STD 58, RFC 2580 [RFC2580].

3. Use Cases

Requirements for power and energy monitoring for networking devices are specified in [EMAN-REQ]. The requirements in [EMAN-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, the scope of the MIB modules in this document is broader than that specified in [EMAN-REQ]. Several use cases for Energy Management have been identified in the "Energy Management (EMAN) Applicability Statement" [EMAN-AS].

4. Terminology

The definitions of basic terms like Power Monitor, Power Monitor Parent, Power Monitor Child, Power Monitor Meter Domain, Power State can be found in the Power Management Architecture [EMAN-FRAMEWORK].

EDITOR'S NOTE: it is foreseen that some more term will follow such a Proxy, Aggregator, Energy Management, etc...

Power State Set

A Power State Set is defined as a sequence of incremental energy saving modes of a device. The elements of this set can be viewed as an interface for the underlying device-implemented power settings of a device. Examples of Power State Sets include DTMF [DMTF], IEEE1621 [IEEE1621], ACPI [ACPI] and EMAN.

Power State

A Power State is defined as a specific power setting for a Power Monitor (e.g., shut, hibernate, sleep, high). Within the context of a Power State Set, the Power State of a device is one of the power saving modes in that Power State Set.

EDITOR'S NOTE: the definitions of Power State Series and Power State should be copied over in [EMAN-FRAMEWORK], and referenced here.

5. Architecture Concepts Applied to the MIB Module

This section describes the concepts specified in the Power Monitor Architecture [EMAN-FRAMEWORK] that pertain to power usage, with specific information related to the MIB module specified in this document. This subsection maps to the section "Architecture High Level Concepts" in the Power Monitoring Architecture [EMAN-FRAMEWORK].

The Energy Monitoring MIB has 2 independent MIB modules. The first MIB module powerMonitorMIB is focused on measurement of power and energy. The second MIB module powerQualityMIB is focused on Power Quality measurement.

The powerMonitorMIB MIB module consists of four tables. The first table pmPowerTable is indexed by pmPowerIndex and pmPowerStateSetIndex. The second table pmPowerStateTable indexed by pmPowerIndex, pmPowerStateSetIndex and pmPowerStateIndex. pmEnergyParametersTable and pmEnergyTable are indexed by pmPowerIndex.

```
pmPowerTable(1)
|
+---pmPowerEntry(1) [pmPowerIndex, pmPowerStateSet]
|
|   +--- --- Integer32          pmPowerIndex(1)
|   +--- --- PowerStateSet    pmPowerStateSet(2)
|   +--- r-n Integer32        pmPower(3)
|   +--- r-n Integer32        pmPowerNamePlate(4)
|   +--- r-n UnitMultiplier   pmPowerUnitMultiplier(5)
|   +--- r-n Integer32        pmPowerAccuracy(6)
|   +--- r-n INTEGER          pmMeasurementCaliber(7)
|   +--- r-n INTEGER          pmPowerCurrentType(8)
|   +--- r-n INTEGER          pmPowerOrigin(9)
|   +--- rwn Integer32        pmPowerAdminState(10)
|   +--- r-n Integer32        pmPowerOperState(11)
|   +--- r-n OwnerString      pmPowerStateEnterReason(12)
|
+---pmPowerStateTable(2)
|   +---pmPowerStateEntry(1)
|       |   [pmPowerIndex,
|       |   pmPowerStateSet,
|       |   pmPowerStateIndex]
|       +--- --- Integer32          pmPowerStateIndex(1)
|       +--- r-n Integer32          pmPowerStateMaxPower (2)
|       +--- r-n UnitMultiplier
|           pmPowerStatePowerUnitMultiplier (3)
|       +--- r-n TimeTicks          pmPowerStateTotalTime(4)
|       +--- r-n Counter64          pmPowerStateEnterCount(5)
|
+pmEnergyParametersTable(1)
+---pmEnergyParametersEntry(1) [pmPowerIndex]
|
|   +--- r-n TimeInterval
|       pmEnergyParametersIntervalLength (1)
|   +--- r-n Integer32
|       pmEnergyParametersIntervalNumber (2)
|   +--- r-n Integer32
|       pmEnergyParametersIntervalMode (3)
```

```

    +--- r-n TimeInterval
        pmEnergyParametersIntervalWindow (4)
    +--- r-n Integer32
        pmEnergyParametersSampleRate (5)
    +--- r-n RowStatus pmEnergyParametersStatus (6)

+pmEnergyTable(1)
+---pmEnergyEntry(1) [pmPowerIndex]

    +--- r-n TimeInterval pmEnergyIntervalStartTime (1)
    +--- r-n Integer32 pmEnergyIntervalEnergyUsed (2)
    +--- r-n UnitMultiplier
        pmEnergyIntervalEnergyUnitMultiplier (3)
    +--- r-n Integer32 pmEnergyIntervalMax (4)
    +--- r-n TimeTicks
        pmEnergyIntervalDiscontinuityTime(5)
    +--- r-n RowStatus pmEnergyParametersStatus (6)

```

The powerQualityMIB consists of four tables. PmACPwrQualityTable is indexed by pmPowerIndex. PmACPwrQualityPhaseTable is indexed by pmPowerIndex and pmPhaseIndex. pmACPwrQualityWyePhaseTable and pmACPwrQualityDelPhaseTable are indexed by pmPowerIndex and pmPhaseIndex.

```

pmPowerTable(1)
+---PmACPwrQualityEntry (1) [pmPowerIndex]
    +----- INTEGER pmACPwrQualityConfiguration (1)
    +--- r-n Integer32 pmACPwrQualityAvgVoltage (2)
    +--- r-n Integer32 pmACPwrQualityAvgCurrent (3)
    +--- r-n Integer32 pmACPwrQualityFrequency (4)
    +--- r-n UnitMultiplier
        pmACPwrQualityPowerUnitMultiplier (5)
    +--- r-n Integer32 pmACPwrQualityPowerAccuracy (6)
    +--- r-n Integer32 pmACPwrQualityTotalActivePower (7)
    +--- r-n Integer32
        pmACPwrQualityTotalReactivePower (8)
    +--- r-n Integer32 pmACPwrQualityTotalApparentPower (9)
    +--- r-n Integer32 pmACPwrQualityTotalPowerFactor(10)
    +--- r-n Integer32 pmACPwrQualityThdAmpheres (11)

+pmACPwrQualityPhaseTable (1)
+---PmACPwrQualityPhaseEntry(1)[pmPowerIndex,
    pmPhaseIndex]

```



```

Internet-Draft          <Energy Monitoring MIB>          July 2011
|
|      +-- r-n Integer32  pmPhaseIndex  (1)
|      +-- r-n Integer32
|      |      pmACPwrQualityPhaseAvgCurrent  (2)
|      +-- r-n Integer32
|      |      pmACPwrQualityPhaseActivePower  (3)
|      +-- r-n Integer32
|      |      pmACPwrQualityPhaseReactivePower  (4)
|      +-- r-n Integer32
|      |      pmACPwrQualityPhaseApparentPower  (5)
|      +-- r-n Integer32
|      |      pmACPwrQualityPhasePowerFactor  (6)
|      +-- r-n Integer32
|      |      pmACPwrQualityPhaseImpedance  (7)
|
+pmACPwrQualityDelPhaseTable (1)
+-- pmACPwrQualityDelPhaseEntry(1)
|
|      [pmPowerIndex,
|      pmPhaseIndex]
|
|      +-- r-n Integer32
|      |      pmACPwrQualityDelPhaseToNextPhaseVoltage  (1)
|      +-- r-n Integer32
|      |      pmACPwrQualityDelThdPhaseToNextPhaseVoltage  (2)
|      +-- r-n Integer32  pmACPwrQualityDelThdCurrent  (3)
|
+pmACPwrQualityWyePhaseTable (1)
+-- pmACPwrQualityWyePhaseEntry (1)
|
|      [pmPowerIndex,
|      pmPhaseIndex]
|
|      +-- r-n Integer32
|      |      pmACPwrQualityWyePhaseToNeutralVoltage  (1)
|      +-- r-n Integer32
|      |      pmACPwrQualityWyePhaseCurrent  (2)
|      +-- r-n Integer32
|      |      pmACPwrQualityWyeThdPhaseToNeutralVoltage  (3)
|
|      .

```

A UML representation of the MIB objects in the two MIB modules are powerMonitorMIB and powerQualityMIB are presented.


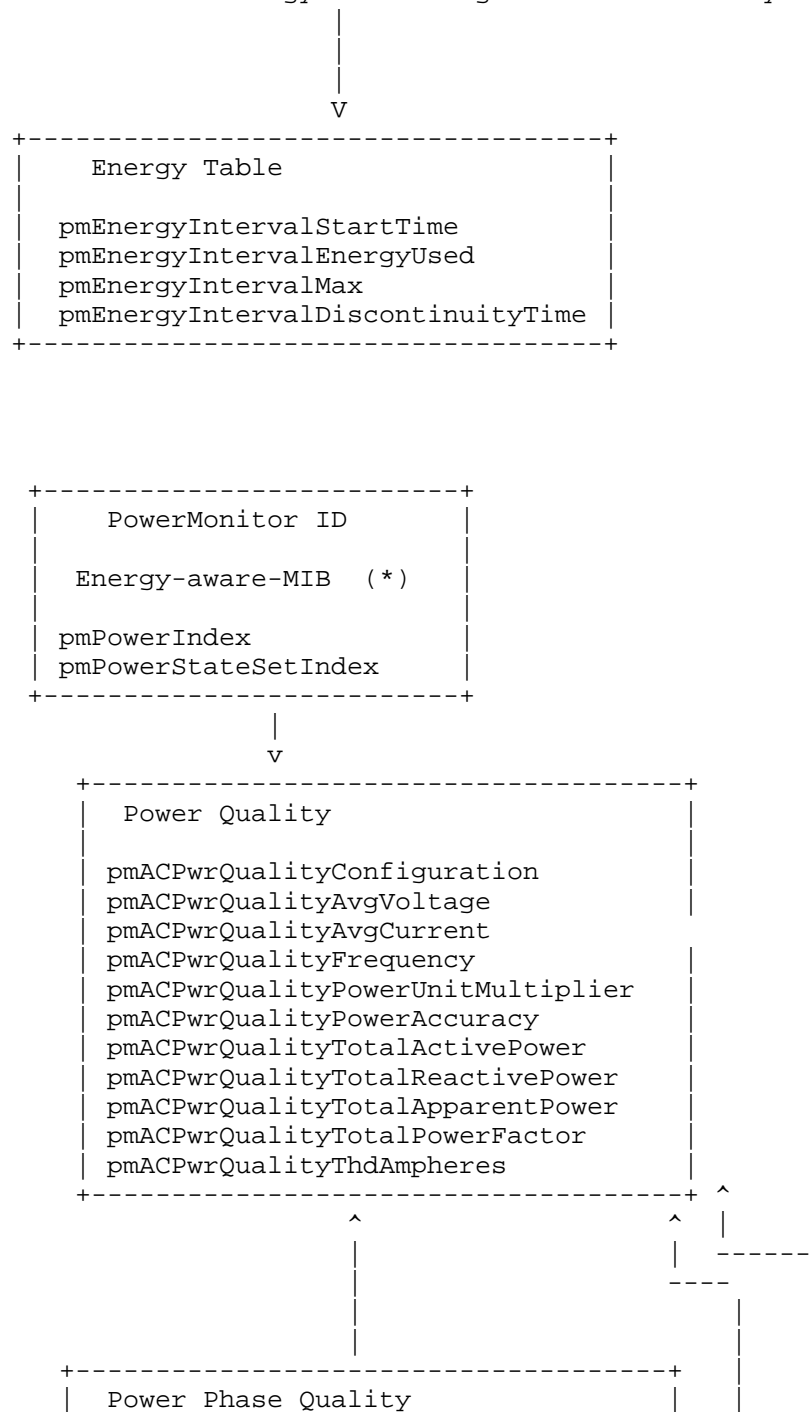


Figure 1:UML diagram for powerMonitor MIB

(*) Link with the ENERGY-AWARE-MIB



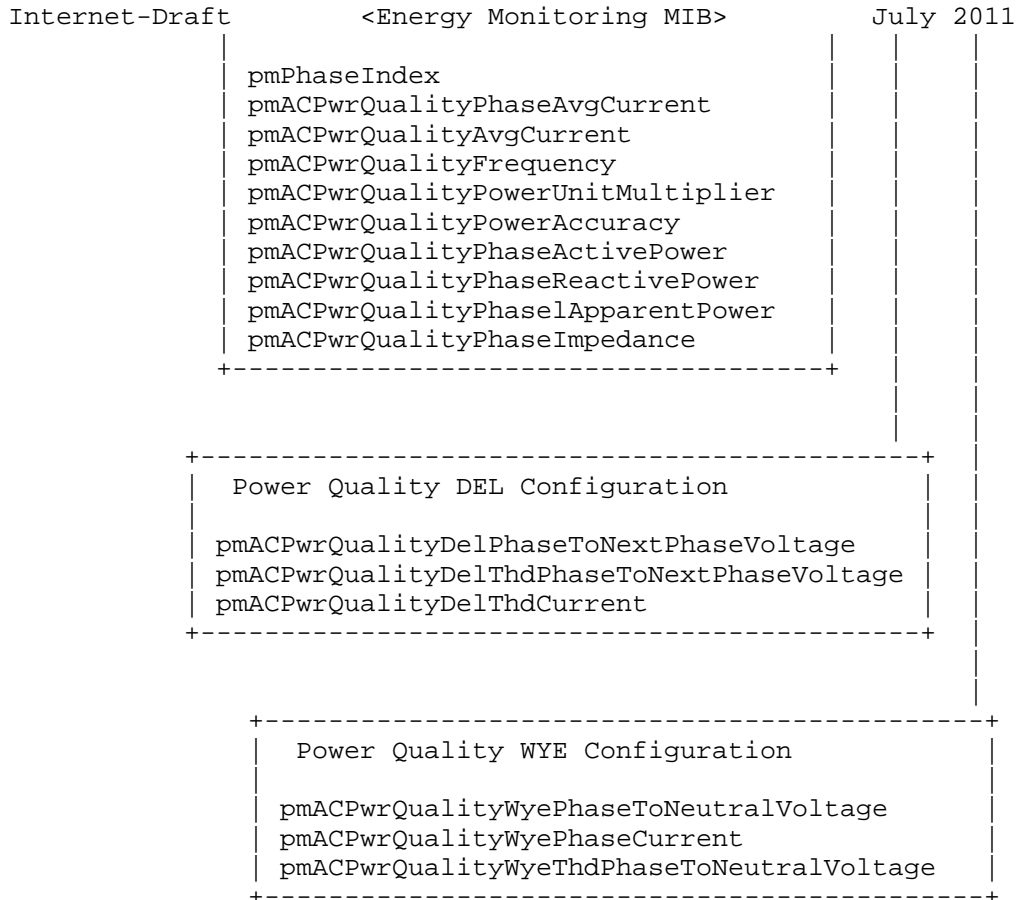


Figure 2: UML diagram for the powerQualityMIB

5.1. Power Monitor Information

Refer to the "Power Monitor Information" section in [EMAN-FRAMEWORK] for background information. An energy aware device is considered an instance of a Power Monitor as defined in the [EMAN-FRAMEWORK].

The Power Monitor identity information is specified in the MIB ENERGY-AWARE-MIB module [EMAN-AWARE-MIB] primary table, i.e. the pmTable. In this table, every Power Monitor SHOULD have a printable name pmName, and MUST HAVE a unique Power Monitor index pmIndex. The ENERGY-AWARE-MIB module returns the relationship (parent/child) between Power Monitors.

5.2. Power State

Refer to the "Power Monitor States" section in [EMAN-FRAMEWORK] for background information.

A Power Monitor may have energy conservation modes called Power States. Between the ON and OFF states of a device, there can be several intermediate energy saving modes. Those energy saving modes are called as Power States.

Power States, which represent universal states of power management of a Power Monitor, are specified by the pmPowerState MIB object. The actual Power State is specified by the pmPowerOperState MIB object, while the pmPowerAdminState MIB object specifies the Power State requested for the Power Monitor. The difference between the values of pmPowerOperState and pmPowerAdminState can be attributed that the Power Monitor is busy transitioning from pmPowerAdminState into the pmPowerOperState, at which point it will update the content of pmPowerOperState. In addition, the possible reason for change in Power State is reported in pmPowerStateEnterReason. Regarding pmPowerStateEnterReason, management stations and Power Monitors should support any format of the owner string dictated by the local policy of the organization. It is suggested that this name contain at least the reason for the transition change, and one or more of the following: IP address, management station name, network manager's name, location, or phone number.

The MIB objects pmPowerOperState, pmPowerAdminState, and pmPowerStateEnterReason are contained in the pmPowerTable MIB table.

The pmPowerStateTable table enumerates the maximum power usage in watts, for every single supported Power State of each Power State Set supported by the Power Monitor. In addition, PowerStateTable provides additional statistics: pmPowerStateEnterCount, the number of times an entity has visited a particular Power State, and pmPowerStateTotalTime, the total time spent in a particular Power State of a Power Monitor.

5.2.1. Power State Set

There are several standards and implementations of Power State Sets. A Power Monitor can support one or multiple Power State Set implementation(s) concurrently.

There are currently three Power State Sets advocated:

```
Reserved(0)
IEEE1621(1) - [IEEE1621]
DMTF(2)      - [DMTF]
EMAN(3)      - [EMAN-MONITORING-MIB]
```

The respective specific states related to each Power State Set are specified in the following sections.

5.2.2. IEEE1621 Power State Set

The IEEE1621 Power State Set [IEEE1621] consists of 3 rudimentary states : on, off or sleep.

on(0) - The device is fully On and all features of the device are in working mode.

off(1) - The device is mechanically switched off and does not consume energy.

sleep(2) - The device is in a power saving mode, and some features may not be available immediately.

5.2.3. DMTF Power State Set

DMTF [DMTF] standards organization has defined a power profile standard based on the CIM (Common Information Model) model that consists of 15 power states ON (2), SleepLight (3), SleepDeep (4), Off-Hard (5), Off-Soft (6), Hibernate(7), PowerCycle Off-Soft (8), PowerCycle Off-Hard (9), MasterBus reset (10), Diagnostic Interrupt (11), Off-Soft-Graceful (12), Off-Hard Graceful (13), MasterBus reset Graceful (14), Power-Cycle Off-Soft Graceful (15), PowerCycle-Hard Graceful (16). DMTF standard is targeted for hosts and computers. Details of the semantics of each Power State within the DMTF Power State Set can be obtained from the DMTF Power State Management Profile specification [DMTF].

DMTF power profile extends ACPI power states. The following table provides a mapping between DMTF and ACPI Power State Set:

DMTF Power State	ACPI Power State
Reserved(0)	
Reserved(1)	
ON (2)	G0-S0
Sleep-Light (3)	G1-S1 G1-S2
Sleep-Deep (4)	G1-S3
Power Cycle (Off-Soft) (5)	G2-S5
Off-hard (6)	G3
Hibernate (Off-Soft) (7)	G1-S4
Off-Soft (8)	G2-S5
Power Cycle (Off-Hard) (9)	G3
Master Bus Reset (10)	G2-S5
Diagnostic Interrupt (11)	G2-S5
Off-Soft Graceful (12)	G2-S5
Off-Hard Graceful (13)	G3
MasterBus Reset Graceful (14)	G2-S5
Power Cycle off-soft Graceful (15)	G2-S5
Power Cycle off-hard Graceful (16)	G3

Figure 3: DMTF and ACPI Powe State Set Mapping

5.2.4. EMAN Power State Set

The EMAN Power State Set represents an attempt for a uniform standard approach to model the different levels of power consumption of a device. The EMAN Power States are an expansion of the basic Power States as defined in IEEE1621 that also

incorporate the Power States defined in ACPI and DMTF. Therefore, in addition to the non-operational states as defined in ACPI and DMTF standards, several intermediate operational states have been defined.

There are twelve Power States, that expand on IEEE1621 on, sleep and off. The expanded list of Power States are divided into six operational states, and six non-operational states. The lowest non-operational state is 1 and the highest is 6. Each non-operational state corresponds to an ACPI state [ACPI] corresponding to Global and System states between G3 (hard-off) and G1 (sleeping). For Each operational state represent a performance state, and may be mapped to ACPI states P0 (maximum performance power) through P5 (minimum performance and minimum power).

An Power Monitor may have fewer Power States than twelve and would then map several policy states to the same power state. Power Monitor with more than twelve states, would choose which twelve to represent as power policy states.

In each of the non-operational states (from mechoff(1) to ready(6)), the Power State preceding it is expected to have a lower power consumption and a longer delay in returning to an operational state:

IEEE1621 Power(off):

mechoff(1) : An off state where no entity features are available. The entity is unavailable. No energy is being consumed and the power connector can be removed. This corresponds to ACPI state G3.

softoff(2) : Similar to mechoff(1), but some components remain powered or receive trace power so that the entity can be awakened from its off state. In softoff(2), no context is saved and the device typically requires a complete boot when awakened. This corresponds to ACPI state G2.

IEEE1621 Power(sleep)

hibernate(3): No entity features are available. The entity may be awakened without requiring a complete boot, but the time for

availability is longer than sleep(4). An example for state hibernate(3) is a save-to-disk state where DRAM context is not maintained. Typically, energy consumption is zero or close to zero. This corresponds to state G1, S4 in ACPI.

- sleep(4) : No entity features are available, except for out-of-band management, for example wake-up mechanisms. The time for availability is longer than standby(5). An example for state sleep(4) is a save-to-RAM state, where DRAM context is maintained. Typically, energy consumption is close to zero. This corresponds to state G1, S3 in ACPI.
- standby(5) : No entity features are available, except for out-of-band management, for example wake-up mechanisms. This mode is analogous to cold-standby. The time for availability is longer than ready(6). For example, the processor context is not maintained. Typically, energy consumption is close to zero. This corresponds to state G1, S2 in ACPI.
- ready(6) : No entity features are available, except for out-of-band management, for example wake-up mechanisms. This mode is analogous to hot-standby. The entity can be quickly transitioned into an operational state. For example, processors are not executing, but processor context is maintained. This corresponds to state G1, S1 in ACPI.

IEEE1621 Power(on):

- lowMinus(7) : Indicates some entity features may not be available and the entity has selected measures/options to provide less than low(8) usage. This corresponds to ACPI State G0. This includes operational states lowMinus(7) to full(12).
- low(8) : Indicates some features may not be available and the entity has taken

measures or selected options to provide less than mediumMinus(9) usage.

mediumMinus(9): Indicates all entity features are available but the entity has taken measures or selected options to provide less than medium(10) usage.

medium(10) : Indicates all entity features are available but the entity has taken measures or selected options to provide less than highMinus(11) usage.

highMinus(11): Indicates all entity features are available and power usage is less than high(12).

high(12) : Indicates all entity features are available and the entity is consuming the highest power.

5.3. Power Monitor Usage Information

Refer to the "Power Monitor Usage Measurement" section in [EMAN-FRAMEWORK] for background information.

For a Power Monitor, power usage is reported using pmPower. The magnitude of measurement is based on the pmPowerUnitMultiplier MIB variable, based on the UnitMultiplier Textual Convention (TC). Power measurement magnitude should conform to the IEC 62053-21 [IEC.62053-21] and IEC 62053-22 [IEC.62053-22] definition of unit multiplier for the SI (System International) units of measure. 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 pmPowerUnitMultiplier. Note that other measurements throughout the two MIB modules in this document use the same mechanism, including pmPowerStatePowerUnitMultiplier, pmEnergyIntervalEnergyUnitMultiplier, and pmACPwrQualityPowerUnitMultiplier.

In addition to knowing the usage and magnitude, it is useful to know how a pmPower measurement was obtained. An NMS can use this to account for the accuracy and nature of the reading

between different implementations. For this `pmPowerOrigin` describes whether the measurements were made at the device itself or from a remote source. The `pmPowerMeasurementCaliber` describes the method that was used to measure the power and can distinguish actual or estimated values. There may be devices in the network, which may not be able to measure or report power consumption. For those devices, the object `pmPowerMeasurementCaliber` shall report that measurement mechanism is "unavailable" and the `pmPower` measurement shall be "0".

The nameplate power rating of a Power Monitor is specified in `pmPowerNameplate` MIB object.

5.4. Optional Power Usage Quality

Refer to the "Optional Power Usage Quality" section in [EMAN-FRAMEWORK] for background information.

The optional `powerQualityMIB` MIB module can be implemented to further describe power usage quality measurement. The `powerQualityMIB` MIB module adheres closely to the IEC 61850 7-2 standard to describe AC measurements.

The `powerQualityMIB` MIB module contains a primary table, the `pmACPwrQualityTable` table, that defines power quality measurements for supported `pmIndex` entities, as a sparse extension of the `pmPowerTable` (with `pmPowerIndex` as primary index). This `pmACPwrQualityTable` table contains such information as the configuration (single phase, DEL 3 phases, WYE 3 phases), voltage, frequency, power accuracy, total active/reactive power/apparent power, amperage, and voltage.

In case of 3-phase power, the `pmACPwrQualityPhaseTable` additional table is populated with power quality measurements per phase (so double indexed by the `pmPowerIndex` and `pmPhaseIndex`). This table, which describes attributes common to both WYE and DEL configurations, contains the average current, active/reactive/apparent power, power factor, and impedance.

In case of 3-phase power with a DEL configuration, the `pmACPwrQualityDelPhaseTable` table describes the phase-to-phase power quality measurements, i.e., voltage and current.

In case of 3-phase power with a Wye configuration, the `pmACPwrQualityWyePhaseTable` table describes the phase-to-neutral power quality measurements, i.e., voltage and current.

5.5. Optional Energy Measurement

Refer to the "Optional Energy and demand Measurement" section in [EMAN-FRAMEWORK] for the definition and terminology information.

It is relevant to measure energy when there are actual power measurements from a Power Monitor, and not when the power measurement is assumed or predicted as specified in the description clause of the object pmPowerMeasurementCaliber.

Two tables are introduced to characterize energy measurement of a Power Monitor: pmEnergyTable and pmEnergyParametersTable. Both energy and demand information can be represented via the pmEnergyTable. Energy information will be an accumulation with no interval. Demand information can be represented as an average accumulation per interval of time.

The pmEnergyParametersTable consists of the parameters defining the duration of measurement intervals in seconds, (pmEnergyParametersIntervalLength), the number of successive intervals to be stored in the pmEnergyTable, (pmEnergyParametersIntervalNumber), the type of measurement technique (pmEnergyParametersIntervalMode), and a sample rate used to calculate the average (pmEnergyParametersSampleRate). Judicious choice of the sampling rate will ensure accurate measurement of energy while not imposing an excessive polling burden.

There are three pmEnergyParametersIntervalMode types used for energy measurement collection: period, sliding, and total. The choices of the the three different modes of collection are based on IEC standard 61850-7-4. Note that multiple pmEnergyParametersIntervalMode types MAY be configured simultaneously.

These three pmEnergyParametersIntervalMode types are illustrated by the following three figures, for which:

- The horizontal axis represents the current time, with the symbol <--- L ---> expressing the pmEnergyParametersIntervalLength, and the pmEnergyIntervalStartTime is represented by S1, S2, S3, S4, ..., Sx where x is the value of pmEnergyParametersIntervalNumber.
- The vertical axis represents the time interval of sampling and the value of pmEnergyIntervalEnergyUsed can be obtained at the

end of the sampling period. The symbol ===== denotes the duration of the sampling period.




Figure 4 : Period pmEnergyParametersIntervalMode

A pmEnergyParametersIntervalMode type of 'period' specifies non-overlapping periodic measurements. Therefore, the next pmEnergyIntervalStartTime is equal to the previous pmEnergyIntervalStartTime plus pmEnergyParametersIntervalLength. $S2=S1+L$; $S3=S2+L$, ...




Figure 5 : Sliding pmEnergyParametersIntervalMode

A pmEnergyParametersIntervalMode type of 'sliding' specifies overlapping periodic measurements.




Figure 4 : Total pmEnergyParametersIntervalMode

A pmEnergyParametersIntervalMode type of 'total' specifies a continuous measurement since the last reset. The value of pmEnergyParametersIntervalNumber should be (1) one and pmEnergyParametersIntervalLength is ignored.

The pmEnergyParametersStatus is used to start and stop energy usage logging. The status of this variable is "active" when all the objects in pmEnergyParametersTable are appropriate which in turn indicates if pmEnergyTable entries exist or not.

The pmEnergyTable consists of energy measurements in pmEnergyIntervalEnergyUsed, the units of the measured energy pmEnergyIntervalEnergyUnitMultiplier, and the maximum observed energy within a window - pmEnergyIntervalMax.

Measurements of the total energy consumed by a Power Monitor may suffer from interruptions in the continuous measurement of energy consumption. In order to indicate such interruptions, the object pmEnergyIntervalDiscontinuityTime is provided for indicating the time of the last interruption of total energy measurement. pmEnergyIntervalDiscontinuityTime shall indicate the sysUpTime [RFC3418] when the device was reset.

The following example illustrates the pmEnergyTable and pmEnergyParametersTable:

First, in order to estimate energy, a time interval to sample energy should be specified, i.e. pmEnergyParametersIntervalLength can be set to "900 seconds" or 15 minutes and the number of consecutive intervals over which

the maximum energy is calculated (pmEnergyParametersIntervalNumber) as "10". The sampling rate internal to the Power Monitor for measurement of power usage (pmEnergyParametersSampleRate) can be "1000 milliseconds", as set by the Power Monitor as a reasonable value. Then, the pmEnergyParametersStatus is set to active (value 1) to indicate that the Power Monitor should start monitoring the usage per the pmEnergyTable.

The indices in the pmEnergyTable are pmPowerIndex, which identifies the Power Monitor, and pmEnergyIntervalStartTime, which denotes the start time of the energy measurement interval based on sysUpTime [RFC3418]. The value of pmEnergyIntervalEnergyUsed is the measured energy consumption over the time interval specified (pmEnergyParametersIntervalLength) based on the Power Monitor internal sampling rate (pmEnergyParametersSampleRate). While choosing the values for the pmEnergyParametersIntervalLength and pmEnergyParametersSampleRate, it is recommended to take into consideration either the network element resources adequate to process and store the sample values, and the mechanism used to calculate the pmEnergyIntervalEnergyUsed. The units are derived from pmEnergyIntervalPowerUnitMultiplier. For example, pmEnergyIntervalPowerUsed can be "100" with pmEnergyIntervalPowerUnits equal to 0, the measured energy consumption of the Power Monitor is 100 watt-hours. The pmEnergyIntervalMax is the maximum energy observed and that can be "150 watt-hours".

The pmEnergyTable has a buffer to retain a certain number of intervals, as defined by pmEnergyParametersIntervalNumber. If the default value of "10" is kept, then the pmEnergyTable contains 10 energymeasurements, including the maximum.

Here is a brief explanation of how the maximum energy can be calculated. The first observed energy measurement value is taken to be the initial maximum. With each subsequent measurement, based on numerical comparison, maximum energy may be updated. The maximum value is retained as long as the measurements are taking place. Based on periodic polling of this table, an NMS could compute the maximum over a longer period, i.e. a month, 3 months, or a year.

5.6. Fault Management

[EMAN-REQ] specifies requirements about Power States such as "the current power state" , "the time of the last state change",

Internet-Draft <Energy Monitoring MIB> July 2011
"the total time spent in each state", "the number of transitions to each state" etc. Some of these requirements are fulfilled explicitly by MIB objects such as pmPowerOperState, pmPowerStateTotalTime and pmPowerStateEnterCount. Some of the other requirements are met via the SNMP NOTIFICATION mechanism. pmPowerStateChange SNMP notification which is generated when the value(s) of pmPowerStateSet, pmPowerOperState, pmPowerAdminState have changed.

6. Discovery

6.1. ENERGY-AWARE-MIB Module Implemented

The NMS must first poll the ENERGY-AWARE-MIB module [EMAN-AWARE-MIB], if available, in order to discover all the Power Monitors and the relationships between those (notion of Parent/Child). In the ENERGY-AWARE-MIB module tables, the Power Monitors are indexed by the pmIndex.

If an implementation of the ENERGY-AWARE-MIB module is available in the local SNMP context, for the same Power Monitor, the pmIndex value (EMAN-AWARE-MIB) MUST be assigned to the pmPowerIndex for The pmPowerIndex characterizes the Power Monitor in the powerMonitorMIB and powerQualityMIB MIB modules (this document).

From there, the NMS must poll the pmPowerStateTable (specified in the powerMonitorMIB module in this document), which enumerates, amongst other things, the maximum power usage. As the entries in pmPowerStateTable table are indexed by the Power Monitor (pmPowerIndex), by the Power State Set (pmPowerStateSetIndex), and by the Power State (pmPowerStateIndex), the maximum power usage is discovered per Power Monitor, per Power State Set, and per Power Usage. In other words, polling the pmPowerStateTable allows the discovery of each Power State within every Power State Set supported by the Power Monitor.

If the Power Monitor is an Aggregator or a Proxy, the MIB module would be populated with the Power Monitor Parent and Children information, which have their own Power Monitor index value (pmPowerIndex). However, the parent/child relationship must be discovered thanks to the ENERGY-AWARE-MIB module.

Finally, the NMS can monitor the Power Quality thanks to the powerQualityMIB MIB module, which reuses the pmPowerIndex to index the Power Monitor.

6.2. ENERGY-AWARE-MIB Module Not Implemented, ENTITY-MIB Implemented

When the ENERGY-AWARE-MIB module [EMAN-AWARE-MIB] is not implemented, the NMS must poll the ENTITY-MIB [RFC4133] in order to discover some more information about the Power Monitors. Indeed, the index for the Power Monitors in the MIB modules specified in this document is the pmPowerIndex, which specifies: "If there is no implementation of the ENERGY-AWARE-MIB module but one of the ENTITY MIB module is available in the local SNMP context, then the same index of an entity MUST be chosen as assigned to the entity by object entPhysicalIndex in the ENTITY MIB module."

As the Section 6.1. , the NMS must then poll the pmPowerStateTable (specified in the powerMonitorMIB module in this document), indexed by the Power Monitor (pmPowerIndex that inherited the entPhysicalIndex value), by the Power State Set (pmPowerStateSetIndex), and by the Power State (pmPowerStateIndex). Then the NMS has discovered every Power State within each Power State Set supported by the Power Monitor.

Note that, without the ENERGY-AWARE-MIB module, the Power Monitor acts as an standalone device, i.e. the notion of parent/child can't be specified.

6.3. ENERGY-AWARE-MIB Module and ENTITY-MIB Not Implemented

If neither the ENERGY-AWARE-MIB module [EMAN-AWARE-MIB] nor of the ENTITY MIB module [RFC4133] are available in the local SNMP context, then this MIB module may choose identity values from a further MIB module providing entity identities.

Note that, without the ENERGY-AWARE-MIB module, the Power Monitor acts as an standalone device, i.e. the notion of parent/child can't be specified.

7. Link with the other IETF MIBs

7.1. Link with the ENTITY MIB and the ENTITY-SENSOR MIB

RFC 4133 [RFC4133] defines the ENTITY MIB module that lists the physical entities of a networking device (router, switch, etc.)

and those physical entities indexed by entPhysicalIndex. From an energy-management standpoint, the physical entities that consume or produce energy are of interest.

RFC 3433 [RFC3433] defines the ENTITY-SENSOR MIB module that provides a standardized way of obtaining information (current value of the sensor, operational status of the sensor, and the data units precision) from sensors embedded in networking devices. Sensors are associated with each index of entPhysicalIndex of the ENTITY MIB [RFC4133]. While the focus of the Power and Energy Monitoring MIB is on measurement of power usage of networking equipment indexed by the ENTITY MIB, this MIB proposes a customized power scale for power measurement and different power state states of networking equipment, and functionality to configure the power state states.

When this MIB module is used to monitor the power usage of devices like routers and switches, the ENTITY MIB and ENTITY-SENSOR MIB SHOULD be implemented. In such cases, the Power Monitors are modeled by the entPhysicalIndex through the pmPhysicalEntity MIB object specified in the pmTable in the ENERGY-AWARE-MIB MIB module [EMAN-AWARE-MIB].

However, the ENTITY-SENSOR MIB [RFC3433] does not have the ANSI C12.x accuracy classes required for electricity (i.e., 1%, 2%, 0.5% accuracy classes). Indeed, entPhySensorPrecision [RFC3433] represents "The number of decimal places of precision in fixed-point sensor values returned by the associated entPhySensorValue object". The ANSI and IEC Standards are used for power measurement and these standards require that we use an accuracy class, not the scientific-number precision model specified in RFC3433. The pmPowerAccuracy MIB object models this accuracy. Note that pmPowerUnitMultiplier represents the scale factor per IEC 62053-21 [IEC.62053-21] and IEC 62053-22 [IEC.62053-22], which is a more logical representation for power measurements (compared to entPhySensorScale), with the mantissa and the exponent values $X * 10^Y$.

Power measurements specifying the qualifier 'UNITS' for each measured value in watts are used in the LLDP-EXT-MED-MIB, POE [RFC3621], and UPS [RFC1628] MIBs. The same 'UNITS' qualifier is used for the power measurement values.

One cannot assume that the ENTITY MIB and ENTITY-SENSOR MIB are implemented for all Power Monitors that need to be monitored. A typical example is a converged building gateway, monitoring several other devices in the building, doing the proxy between SNMP and a protocol like BACNET. Another example is the home

Internet-Draft <Energy Monitoring MIB> July 2011
energy controller. In such cases, the pmPhysicalEntity value contains the zero value, thanks to PhysicalIndexOrZero textual convention.

The pmPowerIndex MIB object has been kept as the unique Power Monitor index. The pmPower is similar to entPhySensorValue [RFC3433] and the pmPowerUnitMultiplier is similar to entPhySensorScale.

7.2. Link with the ENTITY-STATE MIB

For each entity in the ENTITY-MIB [RFC4133], the ENTITY-STATE MIB [RFC4268] specifies the operational states (entStateOper: unknown, enabled, disabled, testing), the alarm (entStateAlarm: unknown, underRepair, critical, major, minor, warning, indeterminate) and the possible values of standby states (entStateStandby: unknown, hotStandby, coldStandby, providingService).

From a power monitoring point of view, in contrast to the entity operational states of entities, Power States are required, as proposed in the Power and Energy Monitoring MIB module. Those Power States can be mapped to the different operational states in the ENTITY-STATE MIB, if a formal mapping is required. For example, the entStateStandby "unknown", "hotStandby", "coldStandby", states could map to the Power State "unknown", "ready", "standby", respectively, while the entStateStandby "providingService" could map to any "low" to "high" Power State.

7.3. Link with the POWER-OVER-ETHERNET MIB

Power-over-Ethernet MIB [RFC3621] provides an energy monitoring and configuration framework for power over Ethernet devices. The RFC introduces a concept of a port group on a switch to define power monitoring and management policy and does not use the entPhysicalIndex as the index. Indeed, the pethMainPseConsumptionPower is indexed by the pethMainPseGroupIndex, which has no mapping with the entPhysicalIndex.

One cannot assume that the Power-over-Ethernet MIB is implemented for all Power Monitors that need to be monitored. A typical example is a converged building gateway, monitoring several other devices in the building, doing the proxy between SNMP and a protocol like BACNET. Another example is the home energy controller. In such cases, the pmethPortIndex and

Internet-Draft <Energy Monitoring MIB> July 2011
pmethPortGrpIndex values contain the zero value, thanks to new
PethPsePortIndexOrZero and textual PethPsePortGroupIndexOrZero
conventions.

However, if the Power-over-Ethernet MIB [RFC3621] is supported,
the Power Monitor pmethPortIndex and pmethPortGrpIndex contain
the pethPsePortIndex and pethPsePortGroupIndex, respectively.

As a consequence, the pmPowerIndex MIB object has been kept as
the unique Power Monitor index.

Note that, even though the Power-over-Ethernet MIB [RFC3621] was
created after the ENTITY-SENSOR MIB [RFC3433], it does not reuse
the precision notion from the ENTITY-SENSOR MIB, i.e. the
entPhySensorPrecision MIB object.

7.4. Link with the UPS MIB

To protect against unexpected power disruption, data centers and
buildings make use of Uninterruptible Power Supplies (UPS). To
protect critical assets, a UPS can be restricted to a particular
subset or domain of the network. UPS usage typically lasts only
for a finite period of time, until normal power supply is
restored. Planning is required to decide on the capacity of the
UPS based on output power and duration of probable power outage.
To properly provision UPS power in a data center or building, it
is important to first understand the total demand required to
support all the entities in the site. This demand can be
assessed and monitored via the Power and Energy Monitoring MIB.

UPS MIB [RFC1628] provides information on the state of the UPS
network. Implementation of the UPS MIB is useful at the
aggregate level of a data center or a building. The MIB module
contains several groups of variables:

- upsIdent: Identifies the UPS entity (name, model, etc.).
- upsBattery group: Indicates the battery state
(upsbatteryStatus, upsEstimatedMinutesRemaining, etc.)
- upsInput group: Characterizes the input load to the UPS
(number of input lines, voltage, current, etc.).
- upsOutput: Characterizes the output from the UPS (number of
output lines, voltage, current, etc.)
- upsAlarms: Indicates the various alarm events.

The measurement of power in the UPS MIB is in Volts, Amperes and Watts. The units of power measurement are RMS volts and RMS Amperes. They are not based on the EntitySensorDataScale and EntitySensorDataPrecision of Entity-Sensor MIB.

Both the Power and Energy Monitoring MIB and the UPS MIB may be implemented on the same UPS SNMP agent, without conflict. In this case, the UPS device itself is the Power Monitor Parent and any of the UPS meters or submeters are the Power Monitor Children.

7.5. Link with the LLDP and LLDP-MED MIBs

The LLDP Protocol is a Data Link Layer protocol used by network devices to advertise their identities, capabilities, and interconnections on a LAN network.

The Media Endpoint Discovery is an enhancement of LLDP, known as LLDP-MED. The LLDP-MED enhancements specifically address voice applications. LLDP-MED covers 6 basic areas: capability discovery, LAN speed and duplex discovery, network policy discovery, location identification discovery, inventory discovery, and power discovery.

Of particular interest to the current MIB module is the power discovery, which allows the endpoint device (such as a PoE phone) to convey power requirements to the switch. In power discovery, LLDP-MED has four Type Length Values (TLVs): power type, power source, power priority and power value. Respectively, those TLVs provide information related to the type of power (power sourcing entity versus powered device), how the device is powered (from the line, from a backup source, from external power source, etc.), the power priority (how important is it that this device has power?), and how much power the device needs.

The power priority specified in the LLDP-MED MIB [LLDP-MED-MIB] actually comes from the Power-over-Ethernet MIB [RFC3621]. If the Power-over-Ethernet MIB [RFC3621] is supported, the exact value from the pethPsePortPowerPriority [RFC3621] is copied over in the lldpXMedRemXPoEPDPowerPriority [LLDP-MED-MIB]; otherwise the value in lldpXMedRemXPoEPDPowerPriority is "unknown". From the Power and Energy Monitoring MIB, it is possible to identify the pethPsePortPowerPriority [RFC3621], thanks to the pmethPortIndex and pmethPortGrpIndex.

The `lldpXMedLocXPoEPDPowerSource` [LLDP-MED-MIB] is similar to `pmPowerOrigin` in indicating if the power for an attached device is local or from a remote device. If the LLDP-MED MIB is supported, the following mapping can be applied to the `pmPowerOrigin`: `lldpXMedLocXPoEPDPowerSource` from `PSE(2)` and `local(3)` can be mapped to `remote(2)` and `self(1)`, respectively.

8. Implementation Scenarios

This section provides an illustrative example scenario for the implementation of the Power Monitor, including Power Monitor Parent and Power Monitor Child relationships.

Example Scenario of a campus network: Switch with PoE Endpoints with further connected Devices

The campus network consists of switches that provide LAN connectivity. The switch with PoE ports is located in wiring closet. PoE IP phones are connected to the switch. The IP phones draw power from the PoE ports of the switch. In addition, a PC is daisy-chained from the IP phone for LAN connectivity.

The IP phone consumes power from the PoE switch, while the PC consumes power from the wall outlet.

The switch has implementations of Entity MIB [RFC4133] and energy-aware MIB [EMAN-AWARE-MIB] while the PC does not have implementation of the Entity MIB, but has an implementation of energy-aware MIB. The switch has the following attributes, `pmPowerIndex` "1", `pmPhysicalEntity` "2", and `pmPowerMonitorId` "UUID 1000". The power usage of the switch is "440 Watts". The switch does not have a Power Monitor Parent.

The PoE switch port has the following attributes: The switch port has `pmPowerIndex` "3", `pmPhysicalEntity` is "12" and `pmPowerMonitorId` is "UUID 1000:3". The power metered at the POE switch port is "12 watts". In this example, the POE switch port has the switch as the Power Monitor Parent, with its `pmParentID` of "1000".

The attributes of the PC are given below. The PC does not have implementation of Entity MIB, and thus does not have `pmPhysicalEntity`. The `pmPowerIndex` (`pmPIndex`) of the PC is "57", the `pmPowerMonitorId` is "UUID 1000:57 ". The PC has a Power Monitor Parent, i.e. the switch port whose

pmPowerMonitorId is "UUID 1000:3". The power usage of the PC is "120 Watts" and is communicated to the switch port.

This example illustrates the important distinction between the Power Monitor Children: The IP phone draws power from the switch, while the PC has LAN connectivity from the phone, but is powered from the wall outlet. However, the Power Monitor Parent sends power control messages to both the Power Monitor Children (IP phone and PC) and the Children react to those messages.



```

Internet-Draft          <Energy Monitoring MIB>          July 2011
=====
| 57          |      0      |  UUID 1000:57 |  UUID 1000:3 | 120          |
=====

```

Figure 1: Example scenario

9. Structure of the MIB

The primary MIB object in this MIB module is the `PowerMonitorMIBObject`. The `pmPowerTable` table of `PowerMonitorMibObject` describes the power measurement attributes of a Power Monitor entity. The notion of identity of the device in terms of uniquely identification of the Power Monitor and its relationship to other entities in the network are addressed in [EMAN-AWARE-MIB].

The power measurement of Power Monitor contains information describing its power usage (`pmPower`) and its current power state (`pmPowerOperState`). In addition to power usage, additional information describing the units of measurement (`pmPowerAccuracy`, `pmPowerUnitMultiplier`), how power usage measurement was obtained (`pmPowerMeasurementCaliber`), the source of power (`pmPowerOrigin`) and the type of power (`pmPowerCurrentTtype`) are described.

A Power Monitor may contain an optional `pmPowerQuality` table that describes the electrical characteristics associated with the current power state and usage.

A Power Monitor may contain an optional `pmEnergyTable` to describe energy measurement information over time.

A Power Monitor may also contain optional battery information associated with this entity.

10. MIB Definitions

```

-- *****
--
--
-- This MIB is used to monitor power usage of network

```


Internet-Draft <Energy Monitoring MIB> July 2011
-- devices
--
-- *****

POWER-MONITOR-MIB DEFINITIONS ::= BEGIN

IMPORTS

 MODULE-IDENTITY,
 OBJECT-TYPE,
 NOTIFICATION-TYPE,
 mib-2,
 Integer32, Counter64, TimeTicks
 FROM SNMPv2-SMI
 TEXTUAL-CONVENTION, DisplayString, RowStatus, TimeInterval
 FROM SNMPv2-TC
 MODULE-COMPLIANCE, NOTIFICATION-GROUP, OBJECT-GROUP
 FROM SNMPv2-CONF
 OwnerString
 FROM RMON-MIB;

powerMonitorMIB MODULE-IDENTITY

 LAST-UPDATED "201107080000Z" -- 8 July 2011
 ORGANIZATION "IETF EMAN Working Group"
 CONTACT-INFO
 "WG charter:
 <http://datatracker.ietf.org/wg/eman/charter/>

 Mailing Lists:
 General Discussion: eman@ietf.org

 To Subscribe:
 <https://www.ietf.org/mailman/listinfo/eman>

 Archive:
 <http://www.ietf.org/mail-archive/web/eman>

 Editors:
 Mouli Chandramouli
 Cisco Systems, Inc.
 Sarjapur Outer Ring Road
 Bangalore,
 IN
 Phone: +91 80 4426 3947
 Email: moulchan@cisco.com

 Brad Schoening
 44 Rivers Edge Drive
 Little Silver, NJ 07739

US

Email: brad@bradschoening.com

Juergen Quittek
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE
Phone: +49 6221 4342-115
Email: quittek@neclab.eu

Thomas Dietz
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
69115 Heidelberg
DE
Phone: +49 6221 4342-128
Email: Thomas.Dietz@nw.neclab.eu

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Degem 1831
Belgium
Phone: +32 2 704 5622
Email: bclaise@cisco.com"

DESCRIPTION

"This MIB is used to monitor power and energy in
devices."

REVISION

"201107080000Z" -- 8 July 2011

DESCRIPTION

"Initial version, published as RFC XXXX."

::= { mib-2 xxx }

powerMonitorMIBNotifs OBJECT IDENTIFIER

::= { powerMonitorMIB 0 }

powerMonitorMIBObjects OBJECT IDENTIFIER

::= { powerMonitorMIB 1 }

Internet-Draft <Energy Monitoring MIB>
powerMonitorMIBConform OBJECT IDENTIFIER
::= { powerMonitorMIB 2 }

July 2011

-- Textual Conventions

PowerStateSet ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"PowerStateSet is a TC that describes the Power State Set a Power Monitor supports. IANA has created a registry of Power State Sets supported by a Power Monitor entity and IANA shall administer the list of Power State Sets.

One byte is used to represent the Power State Set.

field	octets	contents	range
----	-----	-----	-----
1	1	Power State Set	1..255

Note:

the value of Power State Set in network byte order

1 in the first byte indicates IEEE1621 Power State Set

2 in the first byte indicates DMTF Power State Set

3 in the first byte indicates EMAN Power State Set"

REFERENCE

"<http://www.iana.org/assignments/eman>

RFC EDITOR NOTE: please change the previous URL if this is not the correct one after IANA assigned it."

SYNTAX OCTET STRING (SIZE(1))

UnitMultiplier ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

Internet-Draft <Energy Monitoring MIB> July 2011
"The Unit Multiplier is an integer value that represents
the IEEE 61850 Annex A units multiplier associated with
the integer units used to measure the power or energy.

For example, when used with pmPowerUnitMultiplier, -3
represents 10^{-3} or milliwatts."

REFERENCE

"The International System of Units (SI),
National Institute of Standards and Technology,
Spec. Publ. 330, August 1991."

SYNTAX INTEGER {

yocto(-24), -- 10^{-24}
zepto(-21), -- 10^{-21}
atto(-18), -- 10^{-18}
femto(-15), -- 10^{-15}
pico(-12), -- 10^{-12}
nano(-9), -- 10^{-9}
micro(-6), -- 10^{-6}
milli(-3), -- 10^{-3}
units(0), -- 10^0
kilo(3), -- 10^3
mega(6), -- 10^6
giga(9), -- 10^9
tera(12), -- 10^{12}
peta(15), -- 10^{15}
exa(18), -- 10^{18}
zetta(21), -- 10^{21}
yotta(24) -- 10^{24}

}

-- Objects

pmPowerTable OBJECT-TYPE

SYNTAX SEQUENCE OF PmPowerEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
 "This table lists Power Monitors."
 ::= { powerMonitorMIBObjects 1 }

pmPowerEntry OBJECT-TYPE

SYNTAX PmPowerEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
 "An entry describes the power usage of a Power Monitor."

Internet-Draft <Energy Monitoring MIB> July 2011
INDEX { pmPowerIndex, pmPowerStateSetIndex}
 ::= { pmPowerTable 1 }

PmPowerEntry ::= SEQUENCE {
 pmPowerIndex Integer32,
 pmPowerStateSetIndex PowerStateSet,
 pmPower Integer32,
 pmPowerNameplate Integer32,
 pmPowerUnitMultiplier UnitMultiplier,
 pmPowerAccuracy Integer32,
 pmPowerMeasurementCaliber INTEGER,
 pmPowerCurrentType INTEGER,
 pmPowerOrigin INTEGER,
 pmPowerAdminState Integer32,
 pmPowerOperState Integer32,
 pmPowerStateEnterReason OwnerString
}

pmPowerIndex OBJECT-TYPE

SYNTAX Integer32 (0..2147483647)

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"A unique value, for each Power Monitor.

If an implementation of the ENERGY-AWARE-MIB module is available in the local SNMP context, then the same index as the one in the ENERGY-AWARE-MIB MUST be assigned for the identical Power Monitor. In this case, entities without an assigned value for pmIndex cannot be indexed by the pmPowerStateTable.

If there is no implementation of the ENERGY-AWARE-MIB module but one of the ENTITY MIB module is available in the local SNMP context, then the same index of an entity MUST be chosen as assigned to the entity by object entPhysicalIndex in the ENTITY MIB module. In this case, entities without an assigned value for entPhysicalIndex cannot be indexed by the pmPowerStateTable.

If neither the ENERGY-AWARE-MIB module nor of the ENTITY MIB module are available in the local SNMP context, then this MIB module may choose identity values from a further MIB module providing entity identities. In this case the value for each pmPowerIndex must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.

In case that no other MIB modules have been chosen for providing entity identities, Power States can be reported exclusively for the local device on which this table is instantiated. Then this table will have a single entry only and an index value of 0 MUST be used."

::= { pmPowerEntry 1 }

pmPowerStateSetIndex OBJECT-TYPE
SYNTAX PowerStateSet
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
"This object indicates the Power State Set supported by the Power Monitor. The list of Power State Sets and their numbering are administered by IANA"
::= { pmPowerEntry 2 }

pmPower OBJECT-TYPE
SYNTAX Integer32
UNITS "Watts"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
"This object indicates the 'instantaneous' RMS consumption for the Power Monitor. This value is specified in SI units of watts with the magnitude of watts (milliwatts, kilowatts, etc.) indicated separately in pmPowerUnitMultiplier. The accuracy of the measurement is specified in pmPowerAccuracy. The direction of power flow is indicated by the sign on pmPower. If the Power Monitor is consuming power, the pmPower value will be positive. If the Power Monitor is producing power, the pmPower value will be negative.

The pmPower MUST be less than or equal to the maximum power that can be consumed at the power state specified by pmPowerState.

The pmPowerMeasurementCaliber object specifies how the usage value reported by pmPower was obtained. The pmPower value must report 0 if the pmPowerMeasurementCaliber is 'unavailable'. For devices that can not measure or report power, this option can be used."

::= { pmPowerEntry 3 }

pmPowerNameplate OBJECT-TYPE

Internet-Draft	<Energy Monitoring MIB>	July 2011
----------------	-------------------------	-----------

```

SYNTAX      Integer32
UNITS       "Watts"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "This object indicates the rated maximum consumption for
    the fully populated Power Monitor.  The nameplate power
    requirements are the maximum power numbers and, in almost
    all cases, are well above the expected operational
    consumption.  The pmPowerNameplate is widely used for
    power provisioning.  This value is specified in either
    units of watts or voltage and current.  The units are
    therefore SI watts or equivalent Volt-Amperes with the
    magnitude (milliwatts, kilowatts, etc.) indicated
    separately in pmPowerUnitMultiplier."
 ::= { pmPowerEntry 4 }

pmPowerUnitMultiplier OBJECT-TYPE
    SYNTAX      UnitMultiplier
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The magnitude of watts for the usage value in pmPower
        and pmPowerNameplate."
    ::= { pmPowerEntry 5 }

pmPowerAccuracy OBJECT-TYPE
    SYNTAX      Integer32 (0..10000)
    UNITS       "hundredths of percent"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object indicates a percentage value, in 100ths of a
        percent, representing the assumed accuracy of the usage
        reported by pmPower.  For example: The value 1010 means
        the reported usage is accurate to +/- 10.1 percent.  This
        value is zero if the accuracy is unknown or not
        applicable based upon the measurement method.

        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"
    ::= { pmPowerEntry 6 }

pmPowerMeasurementCaliber OBJECT-TYPE
    SYNTAX      INTEGER {

```

```

        unavailable(1) ,
        unknown(2),
        actual(3) ,
        estimated(4),
        presumed(5)
    }
MAX-ACCESS      read-only
STATUS          current
DESCRIPTION
    "This object specifies how the usage value reported by
    pmPower was obtained:

    - unavailable(1): Indicates that the usage is not
    available. In such a case, the pmPower value must be 0
    For devices that can not measure or report power this
    option can be used.

    - unknown(2): Indicates that the way the usage was
    determined is unknown. In some cases, entities report
    aggregate power on behalf of another device. In such
    cases it is not known whether the usage reported is
    actual(2), estimated(3) or presumed (4).

    - actual(3): Indicates that the reported usage was
    measured by the entity through some hardware or direct
    physical means. The usage data reported is not presumed
    (4) or estimated (3) but the real apparent current energy
    consumption rate.

    - estimated(4): Indicates that the usage was not
    determined by physical measurement. The value is a
    derivation based upon the device type, state, and/or
    current utilization using some algorithm or heuristic. It
    is presumed that the entity's state and current
    configuration were used to compute the value.

    - presumed(5): Indicates that the usage was not
    determined by physical measurement, algorithm or
    derivation. The usage was reported based upon external
    tables, specifications, and/or model information. For
    example, a PC Model X draws 200W, while a PC Model Y
    draws 210W"
```

```
 ::= { pmPowerEntry 7 }
```

```

pmPowerCurrentType OBJECT-TYPE
    SYNTAX      INTEGER {
        ac(1),
        dc(2),
```



```
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "This object indicates whether the pmUsage for the Power
        Monitor reports alternative current AC(1), direct current
        DC(2), or that the current type is unknown(3)."
```

::= { pmPowerEntry 8 }

```
pmPowerOrigin OBJECT-TYPE
    SYNTAX INTEGER {
        self (1),
        remote (2)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "This object indicates the source of power measurement
        and can be useful when modeling the power usage of
        attached devices. The power measurement can be performed
        by the entity itself or the power measurement of the
        entity can be reported by another trusted entity using a
        protocol extension. A value of self(1) indicates the
        measurement is performed by the entity, whereas remote(2)
        indicates that the measurement was performed by another
        entity."
```

::= { pmPowerEntry 9 }

```
pmPowerAdminState OBJECT-TYPE
    SYNTAX Integer32 (1..65535)
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "This object specifies the desired Power State for the
        Power Monitor, in the context of the Power State Set
        specified by pmPowerStateSetIndex in this table.
        Possible values of pmPowerAdminState are registered at
        IANA, per Power States Set. A current list of
        assignments can be found at
        <http://www.iana.org/assignments/eman>
        RFC-EDITOR: please check the location after IANA"
```

::= { pmPowerEntry 10 }

```
pmPowerOperState OBJECT-TYPE
    SYNTAX Integer32 (1..65535)
    MAX-ACCESS read-only
```

STATUS

current

DESCRIPTION

"This object specifies the current operational Power State for the Power Monitor, in the context of the Power State Set specified by pmPowerStateSetIndex in this table. Possible values of pmPowerOperState are registered at IANA, per Power States Set. A current list of assignments can be found at <http://www.iana.org/assignments/eman> RFC-EDITOR: please check the list"

```
::= { pmPowerEntry 11 }
```

pmPowerStateEnterReason OBJECT-TYPE

SYNTAX OwnerString

MAX-ACCESS read-create

STATUS current

DESCRIPTION

"This string object describes the reason for the pmPowerAdminState transition Alternatively, this string may contain with the entity that configured this Power Monitor to this Power State."

DEFVAL { "" }

```
::= { pmPowerEntry 12 }
```

pmPowerStateTable OBJECT-TYPE

SYNTAX SEQUENCE OF PmPowerStateEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"This table enumerates the maximum power usage, in watts, for every single supported Power State of each Power Monitor.

This table has an expansion-dependent relationship on the pmPowerTable, containing rows describing each Power State for the corresponding Power Monitor. For every Power Monitor in the pmPowerTable, there is a corresponding entry in this table."

```
::= { powerMonitorMIBObjects 2 }
```

pmPowerStateEntry OBJECT-TYPE

SYNTAX PmPowerStateEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"A pmPowerStateEntry extends a corresponding pmPowerEntry. This entry displays max usage values at

Internet-Draft <Energy Monitoring MIB> July 2011
every single possible Power State supported by the Power
Monitor.
For example, given the values of a Power Monitor
corresponding to a maximum usage of 11W at the
state 1 (mechoff), 6 (ready), 8 (mediumMinus), 12 (High):

State	MaxUsage	Units
1 (mechoff)	0	W
2 (softoff)	0	W
3 (hibernate)	0	W
4 (sleep)	0	W
5 (standby)	0	W
6 (ready)	8	W
7 (lowMinus)	8	W
8 (low)	11	W
9 (medimMinus)	11	W
10 (medium)	11	W
11 (highMinus)	11	W
12 (high)	11	W

Furthermore, this table extends to return the total time
in each Power State, along with the number of times a
particular Power State was entered."

```

INDEX      {
            pmPowerIndex,
            pmPowerStateSetIndex,
            pmPowerStateIndex
        }

 ::= { pmPowerStateTable 1 }

PmPowerStateEntry ::= SEQUENCE {
    pmPowerStateIndex          Integer32,
    pmPowerStateMaxPower       Integer32,
    pmPowerStatePowerUnitMultiplier  UnitMultiplier,
    pmPowerStateTotalTime      TimeTicks,
    pmPowerStateEnterCount     Counter64
}

pmPowerStateIndex OBJECT-TYPE
    SYNTAX          Integer32 (1..65535)
    MAX-ACCESS       not-accessible
    STATUS           current
    DESCRIPTION
        "This object specifies the Power State for the Power
        Monitor, in the context of the Power State Set specified
        by pmPowerStateSetIndex in this table."

```

Internet-Draft <Energy Monitoring MIB> July 2011

 This object specifies the index of the Power State of the Power Monitor within a Power State Set. The semantics of the specific Power State can be obtained from the Power State Set definition."

::= { pmPowerStateEntry 1 }

pmPowerStateMaxPower OBJECT-TYPE

SYNTAX Integer32

UNITS "Watts"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the maximum power for the Power Monitor at the particular Power State. This value is specified in SI units of watts with the magnitude of the units (milliwatts, kilowatts, etc.) indicated separately in pmPowerStatePowerUnitMultiplier. If the maximum power is not known for a certain Power State, then the value is encoded as 0xFFFF.

For Power States not enumerated, the value of pmPowerStateMaxPower might be interpolated by using the next highest supported Power State."

::= { pmPowerStateEntry 3 }

pmPowerStatePowerUnitMultiplier OBJECT-TYPE

SYNTAX UnitMultiplier

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The magnitude of watts for the usage value in pmPowerStateMaxPower."

::= { pmPowerStateEntry 4 }

pmPowerStateTotalTime OBJECT-TYPE

SYNTAX TimeTicks

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the total time in hundreds of seconds that the Power Monitor has been in this power state since the last reset, as specified in the sysUpTime."

::= { pmPowerStateEntry 5 }

pmPowerStateEnterCount OBJECT-TYPE

SYNTAX Counter64

Internet-Draft <Energy Monitoring MIB> July 2011

MAX-ACCESS read-only

STATUS current

DESCRIPTION

 "This object indicates how often the Power Monitor has entered this power state, since the last reset of the device as specified in the sysUpTime."

::= { pmPowerStateEntry 6 }

pmEnergyParametersTable OBJECT-TYPE

SYNTAX SEQUENCE OF PmEnergyParametersEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

 "This table is used to configure the parameters for Energy measurement collection in the table pmEnergyTable."

::= { powerMonitorMIBObjects 4 }

pmEnergyParametersEntry OBJECT-TYPE

SYNTAX PmEnergyParametersEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

 "An entry controls an energy measurement in pmEnergyTable."

INDEX { pmPowerIndex }

::= { pmEnergyParametersTable 1 }

PmEnergyParametersEntry ::= SEQUENCE {

 pmEnergyParametersIntervalLength TimeInterval,

 pmEnergyParametersIntervalNumber Integer32,

 pmEnergyParametersIntervalMode Integer32,

 pmEnergyParametersIntervalWindow TimeInterval,

 pmEnergyParametersSampleRate Integer32,

 pmEnergyParametersStatus RowStatus

}

pmEnergyParametersIntervalLength OBJECT-TYPE

SYNTAX TimeInterval

UNITS "Seconds"

MAX-ACCESS read-create

STATUS current

DESCRIPTION

 "This object indicates the length of time in seconds over which to compute the average pmEnergyIntervalEnergyUsed measurement in the pmEnergyTable table. The computation is based on the Power Monitor's internal sampling rate of

```

Internet-Draft          <Energy Monitoring MIB>          July 2011
    power consumed or produced by the Power Monitor. The
    sampling rate is the rate at which the power monitor can
    read the power usage and may differ based on device
    capabilities. The average energy consumption is then
    computed over the length of the interval."
    DEFVAL { 900 }
    ::= { pmEnergyParametersEntry 1 }

pmEnergyParametersIntervalNumber OBJECT-TYPE
    SYNTAX          Integer32
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "The number of intervals maintained in the pmEnergyTable.
        Each interval is characterized by a specific
        pmEnergyIntervalStartTime, used as an index to the table
        pmEnergyTable . Whenever the maximum number of entries is
        reached, the measurement over the new interval replaces
        the oldest measurement , except if the oldest measurement
        were to be the maximum pmEnergyIntervalMax, in which case
        the measurement the measurement over the next oldest
        interval is replaced."
    DEFVAL { 10 }
    ::= { pmEnergyParametersEntry 2 }

pmEnergyParametersIntervalMode OBJECT-TYPE
    SYNTAX          INTEGER {
                        period(1),
                        sliding(2),
                        total(3)
                    }
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "A control object to define the mode of interval calculation
        for the computation of the average
        pmEnergyIntervalEnergyUsed measurement in the pmEnergyTable
        table.
        A mode of period(1) specifies non-overlapping periodic
        measurements.

        A mode of sliding(2) specifies overlapping sliding windows
        where the interval between the start of one interval and
        the next is defined in pmEnergyParametersIntervalWindow.

        A mode of total(3) specifies non-periodic measurement. In
        this mode only one interval is used as this is a

```

```

Internet-Draft          <Energy Monitoring MIB>          July 2011
    continuous measurement since the last reset. The value of
    pmEnergyParametersIntervalNumber should be (1) one and
    pmEnergyParametersIntervalLength is ignored. "
 ::= { pmEnergyParametersEntry 3 }

pmEnergyParametersIntervalWindow OBJECT-TYPE
    SYNTAX          TimeInterval
    UNITS            "Seconds"
    MAX-ACCESS       read-create
    STATUS           current
    DESCRIPTION
        "The length of the duration window between the starting
        time of one sliding window and the next starting time in
        seconds, in order to compute the average
        pmEnergyIntervalEnergyUsed measurement in the pmEnergyTable
        table. This is valid only when the
        pmEnergyParametersIntervalMode is sliding(2). The
        pmEnergyParametersIntervalWindow value should be a multiple
        of pmEnergyParametersSampleRate."
        ::= { pmEnergyParametersEntry 4 }

pmEnergyParametersSampleRate OBJECT-TYPE
    SYNTAX          Integer32
    UNITS            "Milliseconds"
    MAX-ACCESS       read-create
    STATUS           current
    DESCRIPTION
        "The sampling rate, in milliseconds, at which the Power
        Monitor should poll power usage in order to compute the
        average pmEnergyIntervalEnergyUsed measurement in the
        table pmEnergyTable. The Power Monitor should initially
        set this sampling rate to a reasonable value, i.e., a
        compromise between intervals that will provide good
        accuracy by not being too long, but not so short that
        they affect the Power Monitor performance by requesting
        continuous polling. If the sampling rate is unknown, the
        value 0 is reported. The sampling rate should be selected
        so that pmEnergyParametersIntervalWindow is a multiple of
        pmEnergyParametersSampleRate."
        DEFVAL { 1000 }
        ::= { pmEnergyParametersEntry 5 }

pmEnergyParametersStatus OBJECT-TYPE
    SYNTAX          RowStatus
    MAX-ACCESS       read-create
    STATUS           current
    DESCRIPTION

```

```

Internet-Draft          <Energy Monitoring MIB>          July 2011
    "The status of this row. The pmEnergyParametersStatus is
    used to start or stop energy usage logging. An entry
    status may not be active(1) unless all objects in the
    entry have an appropriate value. If this object is not
    equal to active(1), all associated usage-data logged into
    the pmEnergyTable will be deleted. The data can be
    destroyed by setting up the pmEnergyParametersStatus to
    destroy(2)."
```

```

 ::= { pmEnergyParametersEntry 6 }

pmEnergyTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF PmEnergyIntervalEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This table lists Power Monitor energy measurements.
        Entries in this table are only created if the
        corresponding value of object pmPowerMeasurementCaliber
        is active(2), i.e., if the power is actually metered."
    ::= { powerMonitorMIBObjects 5 }

pmEnergyIntervalEntry OBJECT-TYPE
    SYNTAX          PmEnergyIntervalEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "An entry describing energy measurements."
    INDEX { pmPowerIndex, pmEnergyParametersIntervalMode,
pmEnergyIntervalStartTime }
    ::= { pmEnergyTable 1 }

PmEnergyIntervalEntry ::= SEQUENCE {
    pmEnergyIntervalStartTime      TimeTicks,
    pmEnergyIntervalEnergyUsed     Integer32,
    pmEnergyIntervalEnergyUnitMultiplier UnitMultiplier,
    pmEnergyIntervalMax            Integer32,
    pmEnergyIntervalDiscontinuityTime TimeTicks
}

pmEnergyIntervalStartTime OBJECT-TYPE
    SYNTAX          TimeTicks
    UNITS           "hundredths of seconds"
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION

```



```

Internet-Draft          <Energy Monitoring MIB>          July 2011
    "The time (in hundredths of a second) since the
    network management portion of the system was last
    re-initialized, as specified in the sysUpTime [RFC3418].
    This object is useful for reference of interval periods
    for which the energy is measured."
 ::= { pmEnergyIntervalEntry 1 }

pmEnergyIntervalEnergyUsed OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "Watt-hours"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object indicates the energy used in units of watt-
        hours for the Power Monitor over the defined interval.
        This value is specified in the common billing units of
        watt-hours with the magnitude of watt-hours (kW-Hr, MW-
        Hr, etc.) indicated separately in
        pmEnergyIntervalEnergyUnitMultiplier."
 ::= { pmEnergyIntervalEntry 2 }

pmEnergyIntervalEnergyUnitMultiplier OBJECT-TYPE
    SYNTAX      UnitMultiplier
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object is the magnitude of watt-hours for the
        energy field in pmEnergyIntervalEnergyUsed."
 ::= { pmEnergyIntervalEntry 3 }

pmEnergyIntervalMax OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "Watt-hours"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object is the maximum energy ever observed in
        pmEnergyIntervalEnergyUsed since the monitoring started.
        This value is specified in the common billing units of
        watt-hours with the magnitude of watt-hours (kW-Hr, MW-
        Hr, etc.) indicated separately in
        pmEnergyIntervalEnergyUnits."
 ::= { pmEnergyIntervalEntry 4 }

pmEnergyIntervalDiscontinuityTime OBJECT-TYPE
    SYNTAX      TimeTicks
    MAX-ACCESS  read-only

```

```

Internet-Draft          <Energy Monitoring MIB>          July 2011
STATUS                  current
DESCRIPTION
    "The value of sysUpTime [RFC3418] on the most recent
    occasion at which any one or more of this entity's energy
    consumption counters suffered a discontinuity. If no such
    discontinuities have occurred since the last re-
    initialization of the local management subsystem, then
    this object contains a zero value."
 ::= { pmEnergyIntervalEntry 5 }

-- Notifications

pmPowerStateChange NOTIFICATION-TYPE
    OBJECTS              {pmPowerAdminState, pmPowerOperState,
pmPowerStateEnterReason}
    STATUS                current
    DESCRIPTION
        "The SNMP entity generates the PmPowerStateChange when
        the value(s) of pmPowerAdminState or pmPowerOperState,
        in the context of the Power State Set, have changed for
        the Power Monitor represented by the pmPowerIndex."
 ::= { powerMonitorMIBNotifs 1 }

-- Conformance

powerMonitorMIBCompliances OBJECT IDENTIFIER
 ::= { powerMonitorMIB 3 }

powerMonitorMIBGroups OBJECT IDENTIFIER
 ::= { powerMonitorMIB 4 }

powerMonitorMIBFullCompliance MODULE-COMPLIANCE
STATUS                current
DESCRIPTION
    "When this MIB is implemented with support for
    read-create, then such an implementation can
    claim full compliance. Such devices can then
    be both monitored and configured with this MIB."
MODULE                -- this module
MANDATORY-GROUPS {
    powerMonitorMIBTableGroup,
    powerMonitorMIBStateTableGroup,
    powerMonitorMIBEnergyTableGroup,
    powerMonitorMIBEnergyParametersTableGroup,
    powerMonitorMIBNotifGroup
}
 ::= { powerMonitorMIBCompliances 1 }

```

```

Internet-Draft          <Energy Monitoring MIB>          July 2011
powerMonitorMIBReadOnlyCompliance MODULE-COMPLIANCE
    STATUS              current
    DESCRIPTION
        "When this MIB is implemented without support for
        read-create (i.e. in read-only mode), then such an
        implementation can claim read-only compliance.  Such a
        device can then be monitored but can not be configured
        with this MIB."
    MODULE              -- this module
    MANDATORY-GROUPS {
        powerMonitorMIBTableGroup,
        powerMonitorMIBStateTableGroup,
        powerMonitorMIBNotifGroup
    }

    OBJECT              pmPowerOperState
    MIN-ACCESS          read-only
    DESCRIPTION
        "Write access is not required."
        ::= { powerMonitorMIBCompliances 2 }

-- Units of Conformance

powerMonitorMIBTableGroup OBJECT-GROUP
    OBJECTS
        {
            pmPower,
            pmPowerNameplate,
            pmPowerUnitMultiplier,
            pmPowerAccuracy,
            pmPowerMeasurementCaliber,
            pmPowerCurrentType,
            pmPowerOrigin,
            pmPowerAdminState,
            pmPowerOperState,
            pmPowerStateEnterReason
        }
    STATUS              current
    DESCRIPTION
        "This group contains the collection of all the objects
        related to the PowerMonitor."
        ::= { powerMonitorMIBGroups 1 }

powerMonitorMIBStateTableGroup OBJECT-GROUP
    OBJECTS
        {
            pmPowerStateMaxPower,
            pmPowerStatePowerUnitMultiplier,
            pmPowerStateTotalTime,
            pmPowerStateEnterCount
        }

```

```
    }
    STATUS          current
    DESCRIPTION
        "This group contains the collection of all the
        objects related to the Power State."
    ::= { powerMonitorMIBGroups 2 }
```

```
powerMonitorMIBEnergyParametersTableGroup OBJECT-GROUP
    OBJECTS          {
        pmEnergyParametersIntervalLength,
        pmEnergyParametersIntervalNumber,
        pmEnergyParametersIntervalMode,
        pmEnergyParametersIntervalWindow,
        pmEnergyParametersSampleRate,
        pmEnergyParametersStatus
    }
    STATUS          current
    DESCRIPTION
        "This group contains the collection of all the objects
        related to the configuration of the Energy Table."
    ::= { powerMonitorMIBGroups 3 }
```

```
powerMonitorMIBEnergyTableGroup OBJECT-GROUP
    OBJECTS          {
        -- Note that object
        -- pmEnergyIntervalStartTime is not
        -- included since it is not-accessible

        pmEnergyIntervalEnergyUsed,
        pmEnergyIntervalEnergyUnitMultiplier,
        pmEnergyIntervalMax,
        pmEnergyIntervalDiscontinuityTime
    }
    STATUS          current
    DESCRIPTION
        "This group contains the collection of all the objects
        related to the Energy Table."
    ::= { powerMonitorMIBGroups 4 }
```

```
powerMonitorMIBNotifGroup NOTIFICATION-GROUP
    NOTIFICATIONS    {
        pmPowerStateChange
    }
    STATUS          current
```

```

Internet-Draft          <Energy Monitoring MIB>          July 2011
    DESCRIPTION
        "This group contains the notifications for the power and
        energy monitoring MIB Module."
        ::= { powerMonitorMIBGroups 5 }

END

-- *****
--
-- This MIB module is used to monitor power quality of networked
-- devices with measurements.
--
-- This MIB module is an extension of powerMonitorMIB module.
--
-- *****

POWER-QUALITY-MIB DEFINITIONS ::= BEGIN

IMPORTS
    MODULE-IDENTITY,
    OBJECT-TYPE,
    mib-2,
    Integer32
        FROM SNMPv2-SMI
    MODULE-COMPLIANCE,
    OBJECT-GROUP
        FROM SNMPv2-CONF
    UnitMultiplier, pmPowerIndex
        FROM POWER-MONITOR-MIB
    OwnerString
        FROM RMON-MIB;

powerQualityMIB MODULE-IDENTITY

    LAST-UPDATED      "201107080000Z"          -- 8 July 2011
    ORGANIZATION      "IETF EMAN Working Group"
    CONTACT-INFO
        "WG charter:
        http://datatracker.ietf.org/wg/eman/charter/

        Mailing Lists:
        General Discussion: eman@ietf.org

        To Subscribe:
        https://www.ietf.org/mailman/listinfo/eman

```

Internet-Draft

<Energy Monitoring MIB>

July 2011

Archive:

<http://www.ietf.org/mail-archive/web/eman>

Editors:

Mouli Chandramouli
Cisco Systems, Inc.
Sarjapur Outer Ring Road
Bangalore,
IN
Phone: +91 80 4426 3947
Email: moulchan@cisco.com

Brad Schoening
44 Rivers Edge Drive
Little Silver, NJ 07739
US
Email: brad@bradschoening.com

Juergen Quittek
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE
Phone: +49 6221 4342-115
Email: quittek@neclab.eu

Thomas Dietz
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
69115 Heidelberg
DE
Phone: +49 6221 4342-128
Email: Thomas.Dietz@nw.neclab.eu

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Degem 1831
Belgium
Phone: +32 2 704 5622
Email: bclaise@cisco.com

DESCRIPTION

<Claise, et. Al>

Expires January 8, 2012

[Page 53]

Internet-Draft <Energy Monitoring MIB> July 2011
 "This MIB is used to report AC power quality in
 devices. The table is a sparse augmentation of the
 pmPowerTable table from the powerMonitorMIB module.
 Both three-phase and single-phase power
 configurations are supported."

REVISION

 "201107080000Z" -- 8 July 2011

DESCRIPTION

 "Initial version, published as RFC YYY."

::= { mib-2 yyy }

powerQualityMIBConform OBJECT IDENTIFIER

::= { powerQualityMIB 0 }

powerQualityMIBObjects OBJECT IDENTIFIER

::= { powerQualityMIB 1 }

-- Objects

pmACPwrQualityTable OBJECT-TYPE

SYNTAX SEQUENCE OF PmACPwrQualityEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

 "This table defines power quality measurements for
 supported pmPowerIndex entities. It is a sparse
 extension of the pmPowerTable."

::= { powerQualityMIBObjects 1 }

pmACPwrQualityEntry OBJECT-TYPE

SYNTAX PmACPwrQualityEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

 "This is a sparse extension of the pmPowerTable with
 entries for power quality measurements or
 configuration. Each measured value corresponds to an
 attribute in IEC 61850-7-4 for non-phase measurements
 within the object MMUX."

INDEX { pmPowerIndex }

::= { pmACPwrQualityTable 1 }

```

PmACPwrQualityEntry ::= SEQUENCE {
    pmACPwrQualityConfiguration      INTEGER,
    pmACPwrQualityAvgVoltage         Integer32,
    pmACPwrQualityAvgCurrent         Integer32,
    pmACPwrQualityFrequency          Integer32,
    pmACPwrQualityPowerUnitMultiplier UnitMultiplier,
    pmACPwrQualityPowerAccuracy      Integer32,
    pmACPwrQualityTotalActivePower   Integer32,
    pmACPwrQualityTotalReactivePower Integer32,
    pmACPwrQualityTotalApparentPower Integer32,
    pmACPwrQualityTotalPowerFactor   Integer32,
    pmACPwrQualityThdAmperes         Integer32,
    pmACPwrQualityThdVoltage         Integer32
}

```

pmACPwrQualityConfiguration OBJECT-TYPE

```

SYNTAX INTEGER {
    sngl(1),
    del(2),
    wye(3)
}

```

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"Configuration describes the physical configurations of the power supply lines:

- * alternating current, single phase (SNGL)
- * alternating current, three phase delta (DEL)
- * alternating current, three phase Y (WYE)

Three-phase configurations can be either connected in a triangular delta (DEL) or star Y (WYE) system. WYE systems have a shared neutral voltage, while DEL systems do not. Each phase is offset 120 degrees to each other."

```
 ::= { pmACPwrQualityEntry 1 }
```

pmACPwrQualityAvgVoltage OBJECT-TYPE

```

SYNTAX      Integer32
UNITS       "0.1 Volt AC"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION

```

"A measured value for average 'instantaneous' RMS line voltage. For a 3-phase system, this is the average voltage (V1+V2+V3)/3. IEC 61850-7-4 measured value attribute 'Vol'"

pmACPwrQualityAvgCurrent OBJECT-TYPE
SYNTAX Integer32
UNITS "Amperes"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "A measured value of the current per phase. IEC 61850-
 7-4 attribute 'Amp'."
::= { pmACPwrQualityEntry 3 }

pmACPwrQualityFrequency OBJECT-TYPE
SYNTAX Integer32 (4500..6500) -- UNITS 0.01 Hertz
UNITS "hertz"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "A measured value for the basic frequency of the AC
 circuit. IEC 61850-7-4 attribute 'Hz'."
::= { pmACPwrQualityEntry 4 }

pmACPwrQualityPowerUnitMultiplier OBJECT-TYPE
SYNTAX UnitMultiplier
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "The magnitude of watts for the usage value in
 pmACPwrQualityTotalActivePower,
 pmACPwrQualityTotalReactivePower
 and pmACPwrQualityTotalApparentPower measurements. For
 3-phase power systems, this will also include
 pmACPwrQualityPhaseActivePower,
 pmACPwrQualityPhaseReactivePower and
 pmACPwrQualityPhaseApparentPower"
::= { pmACPwrQualityEntry 5 }

pmACPwrQualityPowerAccuracy OBJECT-TYPE
SYNTAX Integer32 (0..10000)
UNITS "hundredths of percent"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "This object indicates a percentage value, in 100ths of
 a percent, representing the presumed accuracy of
 active, reactive, and apparent power usage reporting.
 For example: 1010 means the reported usage is accurate

Internet-Draft <Energy Monitoring MIB> July 2011
to +/- 10.1 percent. This value is zero if the
accuracy is unknown.

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"
::= { pmACPwrQualityEntry 6 }

pmACPwrQualityTotalActivePower OBJECT-TYPE
SYNTAX Integer32
UNITS "RMS watts"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "A measured value of the actual power delivered to or
 consumed by the load. IEC 61850-7-4 attribute 'TotW'."
::= { pmACPwrQualityEntry 7 }

pmACPwrQualityTotalReactivePower OBJECT-TYPE
SYNTAX Integer32
UNITS "volt-amperes reactive"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "A measured value of the reactive portion of the
 apparent power. IEC 61850-7-4 attribute 'TotVAR'."
::= { pmACPwrQualityEntry 8 }

pmACPwrQualityTotalApparentPower OBJECT-TYPE
SYNTAX Integer32
UNITS "volt-amperes"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "A measured value of the voltage and current which
 determines the apparent power. The apparent power is
 the vector sum of real and reactive power.

 Note: watts and volt-amperes are equivalent units and
 may be combined. IEC 61850-7-4 attribute 'TotVA'."
::= { pmACPwrQualityEntry 9 }

pmACPwrQualityTotalPowerFactor OBJECT-TYPE
SYNTAX Integer32 (-10000..10000)
UNITS "hundredths of percent"
MAX-ACCESS read-only
STATUS current

DESCRIPTION

"A measured value ratio of the real power flowing to the load versus the apparent power. It is dimensionless and expressed here as a percentage value in 100ths of a percent. A power factor of 100% indicates there is no inductance load and thus no reactive power. Power Factor can be positive or negative, where the sign should be in lead/lag (IEEE) form. IEC 61850-7-4 attribute 'TotPF'."

::= { pmACPwrQualityEntry 10 }

pmACPwrQualityThdAmperes OBJECT-TYPE

SYNTAX Integer32 (0..10000)
UNITS "hundredths of percent"
MAX-ACCESS read-only
STATUS current

DESCRIPTION

"A calculated value for the current total harmonic distortion (THD). Method of calculation is not specified. IEC 61850-7-4 attribute 'ThdAmp'."

::= { pmACPwrQualityEntry 11 }

pmACPwrQualityThdVoltage OBJECT-TYPE

SYNTAX Integer32 (0..10000)
UNITS "hundredths of percent"
MAX-ACCESS read-only
STATUS current

DESCRIPTION

"A calculated value for the voltage total harmonic distortion (THD). Method of calculation is not specified. IEC 61850-7-4 attribute 'ThdVol'."

::= { pmACPwrQualityEntry 12 }

pmACPwrQualityPhaseTable OBJECT-TYPE

SYNTAX SEQUENCE OF PmACPwrQualityPhaseEntry
MAX-ACCESS not-accessible
STATUS current

DESCRIPTION

"This table describes 3-phase power quality measurements. It is a sparse extension of the pmACPwrQualityTable."

::= { powerQualityMIBObjects 2 }

pmACPwrQualityPhaseEntry OBJECT-TYPE

SYNTAX PmACPwrQualityPhaseEntry
MAX-ACCESS not-accessible
STATUS current

DESCRIPTION

"An entry describes common 3-phase power quality measurements.

This optional table describes 3-phase power quality measurements, with three entries for each supported pmPowerIndex entity. Entities having single phase power shall not have any entities.

This table describes attributes common to both WYE and DEL. Entities having single phase power shall not have any entries here. It is a sparse extension of the pmACPwrQualityTable.

These attributes correspond to IEC 61850-7.4 MMXU phase measurements."

```
INDEX { pmPowerIndex, pmPhaseIndex }
 ::= { pmACPwrQualityPhaseTable 1 }
```

```
PmACPwrQualityPhaseEntry ::= SEQUENCE {
    pmPhaseIndex                Integer32,
    pmACPwrQualityPhaseAvgCurrent Integer32,
    pmACPwrQualityPhaseActivePower Integer32,
    pmACPwrQualityPhaseReactivePower Integer32,
    pmACPwrQualityPhaseApparentPower Integer32,
    pmACPwrQualityPhasePowerFactor Integer32,
    pmACPwrQualityPhaseImpedance Integer32
}
```

pmPhaseIndex OBJECT-TYPE

```
SYNTAX      Integer32 (0..359)
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "A phase angle typically corresponding to 0, 120, 240."
 ::= { pmACPwrQualityPhaseEntry 1 }
```

pmACPwrQualityPhaseAvgCurrent OBJECT-TYPE

```
SYNTAX      Integer32
UNITS       "Amperes"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "A measured value of the current per phase. IEC 61850-
    7-4 attribute 'A'"
 ::= { pmACPwrQualityPhaseEntry 2 }
```

pmACPwrQualityPhaseActivePower OBJECT-TYPE

```
SYNTAX      Integer32
```

Internet-Draft	<Energy Monitoring MIB>	July 2011
----------------	-------------------------	-----------

```

    UNITS          "RMS watts"
    MAX-ACCESS     read-only
    STATUS         current
    DESCRIPTION
        "A measured value of the actual power delivered to or
        consumed by the load. IEC 61850-7-4 attribute 'W'"
    ::= { pmACPwrQualityPhaseEntry 3 }

pmACPwrQualityPhaseReactivePower OBJECT-TYPE
    SYNTAX         Integer32
    UNITS          "volt-amperes reactive"
    MAX-ACCESS     read-only
    STATUS         current
    DESCRIPTION
        "A measured value of the reactive portion of the
        apparent power. IEC 61850-7-4 attribute 'VAr'"
    ::= { pmACPwrQualityPhaseEntry 4 }

pmACPwrQualityPhaseApparentPower OBJECT-TYPE
    SYNTAX         Integer32
    UNITS          "volt-amperes"
    MAX-ACCESS     read-only
    STATUS         current
    DESCRIPTION
        "A measured value of the voltage and current determines
        the apparent power. Active plus reactive power equals
        the total apparent power.

        Note: Watts and volt-amperes are equivalent units and
        may be combined. IEC 61850-7-4 attribute 'VA'."
    ::= { pmACPwrQualityPhaseEntry 5 }

pmACPwrQualityPhasePowerFactor OBJECT-TYPE
    SYNTAX         Integer32 (-10000..10000)
    UNITS          "hundredths of percent"
    MAX-ACCESS     read-only
    STATUS         current
    DESCRIPTION
        "A measured value ratio of the real power flowing to
        the load versus the apparent power for this phase. IEC
        61850-7-4 attribute 'PF'. Power Factor can be positive
        or negative where the sign should be in lead/lag (IEEE)
        form."
    ::= { pmACPwrQualityPhaseEntry 6 }

pmACPwrQualityPhaseImpedance OBJECT-TYPE
    SYNTAX         Integer32
    UNITS          "volt-amperes"

```

```

Internet-Draft          <Energy Monitoring MIB>          July 2011
    MAX-ACCESS          read-only
    STATUS              current
    DESCRIPTION
    "A measured value of the impedance.  IEC 61850-7-4 attribute
    'Z'."
    ::= { pmACPwrQualityPhaseEntry 7 }

pmACPwrQualityDelPhaseTable OBJECT-TYPE
    SYNTAX              SEQUENCE OF PmACPwrQualityDelPhaseEntry
    MAX-ACCESS          not-accessible
    STATUS              current
    DESCRIPTION
        "This table describes DEL configuration phase-to-phase
        power quality measurements.  This is a sparse extension
        of the pmACPwrQualityPhaseTable."
    ::= { powerQualityMIBObjects 3 }

pmACPwrQualityDelPhaseEntry OBJECT-TYPE
    SYNTAX              PmACPwrQualityDelPhaseEntry
    MAX-ACCESS          not-accessible
    STATUS              current
    DESCRIPTION
        "An entry describes quality attributes of a phase in a
        DEL 3-phase power system.  Voltage measurements are
        provided both relative to each other and zero.

        Measured values are from IEC 61850-7-2 MMUX and THD from
        MHAI objects.

        For phase-to-phase measurements, the pmPhaseIndex is
        compared against the following phase at +120 degrees.
        Thus, the possible values are:

                pmPhaseIndex          Next Phase Angle
                    0                  120
                   120                 240
                   240                  0

        "
    INDEX { pmPowerIndex, pmPhaseIndex}
    ::= { pmACPwrQualityDelPhaseTable 1}

PmACPwrQualityDelPhaseEntry ::= SEQUENCE {
    pmACPwrQualityDelPhaseToNextPhaseVoltage      Integer32,
    pmACPwrQualityDelThdPhaseToNextPhaseVoltage   Integer32,
    pmACPwrQualityDelThdCurrent                    Integer32
}

pmACPwrQualityDelPhaseToNextPhaseVoltage OBJECT-TYPE

```

Internet-Draft	<Energy Monitoring MIB>	July 2011
----------------	-------------------------	-----------

```

SYNTAX      Integer32
UNITS       "0.1 Volt AC"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "A measured value of phase to next phase voltages, where
    the next phase is IEC 61850-7-4 attribute 'PPV'."
 ::= { pmACPwrQualityDelPhaseEntry 2 }

pmACPwrQualityDelThdPhaseToNextPhaseVoltage OBJECT-TYPE
SYNTAX      Integer32 (0..10000)
UNITS       "hundredths of percent"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "A calculated value for the voltage total harmonic
    distortion for phase to next phase. Method of calculation
    is not specified. IEC 61850-7-4 attribute 'ThdPPV'."
 ::= { pmACPwrQualityDelPhaseEntry 3 }

pmACPwrQualityDelThdCurrent OBJECT-TYPE
SYNTAX      Integer32 (0..10000)
UNITS       "hundredths of percent"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "A calculated value for the voltage total harmonic
    distortion (THD) for phase to phase. Method of
    calculation is not specified.
    IEC 61850-7-4 attribute 'ThdPPV'."
 ::= { pmACPwrQualityDelPhaseEntry 4 }

pmACPwrQualityWyePhaseTable OBJECT-TYPE
SYNTAX      SEQUENCE OF PmACPwrQualityWyePhaseEntry
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "This table describes WYE configuration phase-to-neutral
    power quality measurements. This is a sparse extension
    of the pmACPwrQualityPhaseTable."
 ::= { powerQualityMIBObjects 4 }

pmACPwrQualityWyePhaseEntry OBJECT-TYPE
SYNTAX      PmACPwrQualityWyePhaseEntry
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION

```

"This table describes measurements of WYE configuration with phase to neutral power quality attributes. Three entries are required for each supported pmPowerIndex entry. Voltage measurements are relative to neutral.

This is a sparse extension of the pmACPwrQualityPhaseTable.

Each entry describes quality attributes of one phase of a WYE 3-phase power system.

Measured values are from IEC 61850-7-2 MMUX and THD from MHAI objects."

```
INDEX { pmPowerIndex, pmPhaseIndex }  
 ::= { pmACPwrQualityWyePhaseTable 1 }
```

```
PmACPwrQualityWyePhaseEntry ::= SEQUENCE {  
    pmACPwrQualityWyePhaseToNeutralVoltage      Integer32,  
    pmACPwrQualityWyePhaseCurrent                Integer32,  
    pmACPwrQualityWyeThdPhaseToNeutralVoltage    Integer32  
}
```

pmACPwrQualityWyePhaseToNeutralVoltage OBJECT-TYPE

```
SYNTAX      Integer32  
UNITS       "0.1 Volt AC"  
MAX-ACCESS  read-only  
STATUS      current  
DESCRIPTION  
    "A measured value of phase to neutral voltage. IEC  
    61850-7-4 attribute 'PhV'."  
 ::= { pmACPwrQualityWyePhaseEntry 1 }
```

pmACPwrQualityWyePhaseCurrent OBJECT-TYPE

```
SYNTAX      Integer32  
UNITS       "0.1 amperes AC"  
MAX-ACCESS  read-only  
STATUS      current  
DESCRIPTION  
    "A measured value of phase currents. IEC 61850-7-4  
    attribute 'A'."  
 ::= { pmACPwrQualityWyePhaseEntry 2 }
```

pmACPwrQualityWyeThdPhaseToNeutralVoltage OBJECT-TYPE

```
SYNTAX      Integer32 (0..10000)  
UNITS       "hundredths of percent"  
MAX-ACCESS  read-only  
STATUS      current  
DESCRIPTION
```



```

Internet-Draft          <Energy Monitoring MIB>          July 2011
    "A calculated value of the voltage total harmonic
    distortion (THD) for phase to neutral. IEC 61850-7-4
    attribute 'ThdPhV'."
    ::= { pmACPwrQualityWyePhaseEntry 3 }

-- Conformance

powerQualityMIBCompliances OBJECT IDENTIFIER
    ::= { powerQualityMIB 2 }

powerQualityMIBGroups OBJECT IDENTIFIER
    ::= { powerQualityMIB 3 }

powerQualityMIBFullCompliance MODULE-COMPLIANCE
    STATUS          current
    DESCRIPTION
        "When this MIB is implemented with support for read-
        create, then such an implementation can claim full
        compliance. Such devices can then be both monitored and
        configured with this MIB."
    MODULE          -- this module
    MANDATORY-GROUPS {
        powerACPwrQualityMIBTableGroup,
        powerACPwrQualityPhaseMIBTableGroup
    }

    GROUP          powerACPwrQualityDelPhaseMIBTableGroup
    DESCRIPTION
        "This group must only be implemented for a DEL phase
        configuration."

    GROUP          powerACPwrQualityWyePhaseMIBTableGroup
    DESCRIPTION
        "This group must only be implemented for a WYE phase
        configuration."
    ::= { powerQualityMIBCompliances 1 }

-- Units of Conformance

powerACPwrQualityMIBTableGroup OBJECT-GROUP
    OBJECTS          {
        -- Note that object pmPowerIndex is NOT
        -- included since it is not-accessible
        pmACPwrQualityConfiguration,
        pmACPwrQualityAvgVoltage,
        pmACPwrQualityAvgCurrent,
        pmACPwrQualityFrequency,

```

```

    pmACPwrQualityPowerUnitMultiplier,
    pmACPwrQualityPowerAccuracy,
    pmACPwrQualityTotalActivePower,
    pmACPwrQualityTotalReactivePower,
    pmACPwrQualityTotalApparentPower,
    pmACPwrQualityTotalPowerFactor,
    pmACPwrQualityThdAmperes,
    pmACPwrQualityThdVoltage
  }      STATUS      current

DESCRIPTION
  "This group contains the collection of all the power
  quality objects related to the Power Monitor."
 ::= { powerQualityMIBGroups 1 }

```

```

powerACPwrQualityPhaseMIBTableGroup OBJECT-GROUP
OBJECTS
    {
        -- Note that object pmPowerIndex is NOT
        -- included since it is not-accessible
        pmACPwrQualityPhaseAvgCurrent,
        pmACPwrQualityPhaseActivePower,
        pmACPwrQualityPhaseReactivePower,
        pmACPwrQualityPhaseApparentPower,
        pmACPwrQualityPhasePowerFactor,
        pmACPwrQualityPhaseImpedance
    }
STATUS      current
DESCRIPTION
  "This group contains the collection of all 3-phase power
  quality objects related to the Power State."
 ::= { powerQualityMIBGroups 2 }

```

```

powerACPwrQualityDelPhaseMIBTableGroup OBJECT-GROUP
OBJECTS
    {
        -- Note that object pmPowerIndex and
        -- pmPhaseIndex are NOT included
        -- since they are not-accessible
        pmACPwrQualityDelPhaseToNextPhaseVoltage ,
        pmACPwrQualityDelThdPhaseToNextPhaseVoltage,
        pmACPwrQualityDelThdCurrent
    }
STATUS      current
DESCRIPTION
  "This group contains the collection of all quality
  attributes of a phase in a DEL 3-phase power system."
 ::= { powerQualityMIBGroups 3 }

```

```

powerACPwrQualityWyePhaseMIBTableGroup OBJECT-GROUP

```

```

Internet-Draft      <Energy Monitoring MIB>      July 2011
OBJECTS      {
    -- Note that object pmPowerIndex and
    -- pmPhaseIndex are NOT included
    -- since they are not-accessible
    pmACPwrQualityWyePhaseToNeutralVoltage,
    pmACPwrQualityWyePhaseCurrent,
    pmACPwrQualityWyeThdPhaseToNeutralVoltage
}
STATUS      current
DESCRIPTION
    "This group contains the collection of all WYE
    configuration phase-to-neutral power quality
    measurements."
 ::= { powerQualityMIBGroups 4 }

END

```

11. Security Considerations

Some of the readable objects in these 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 even GET and/or NOTIFY access to these objects and possibly to even encrypt the values of these objects when sending them over the network via SNMP.

There are a number of management objects defined in these MIB modules with a MAX-ACCESS clause of read-write and/or read-create. Such objects MAY be considered sensitive or vulnerable in some network environments. The support for SET operations in a non-secure environment without proper protection can have a negative effect on network operations. The following are the tables and objects and their sensitivity/vulnerability:

- Unauthorized changes to the pmPowerOperState (via the pmPowerAdminState) MAY disrupt the power settings of the different Power Monitors, and therefore the state of functionality of the respective Power Monitors.
- Unauthorized changes to the pmEnergyParametersTable MAY disrupt energy measurement in the pmEnergyTable table.

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. IANA Considerations

12.1. IANA Considerations for the MIB Modules

The MIB modules in this document uses the following IANA-assigned OBJECT IDENTIFIER values recorded in the SMI Numbers registry:

Descriptor	OBJECT IDENTIFIER value
-----	-----
PowerMonitorMIB	{ mib-2 xxx }
powerQualityMIB	{ mib-2 yyy }

Additions to the MIB modules are subject to Expert Review [RFC5226], i.e., review by one of a group of experts designated by an IETF Area Director. The group of experts MUST check the requested MIB objects for completeness and accuracy of the description. Requests for MIB objects that duplicate the functionality of existing objects SHOULD be declined. The smallest available OIDs SHOULD be assigned to the new MIB objects. The specification of new MIB objects SHOULD follow the structure specified in Section 10. and MUST be published using a well-established and persistent publication medium.

12.2. IANA Registration of new Power State Set

This document specifies an initial set of Power State Sets. The list of these Power State Sets with their numeric identifiers is given in Section 5.2.1. The Internet Assigned Numbers Authority (IANA) has created a new registry for Power State Sets numeric

Internet-Draft <Energy Monitoring MIB> July 2011
identifiers and filled it with the initial list as in Section
5.2.1. New Assignments to Power State Sets shall be
administered by IANA and the guidelines and procedures are
listed in this Section.

New assignments in Power State Sets require a Standards Action [RFC5226], i.e., they are to be made via Standards Track RFCs approved by the IESG. The new Power State Set based on the following guidelines; firstly check if there are devices or entities that have implementations of the proposed Power State Set or secondly, if the new Power State Set has been adopted or approved by the respective energy management standards organizations. A pure vendor specific implementation of Power State Set shall not be adopted; since it would lead to proliferation of Power State Sets.

12.2.1. IANA Registration of the IEEE1621 Power State Set

This document specifies a set of values for the IEEE1621 Power State Set [IEEE1621]. The list of these values with their identifiers is given in Section 5.2.1. The Internet Assigned Numbers Authority (IANA) created a new registry for IEEE1621 Power State Set identifiers and filled it with the initial list in Section 5.2.2.

New assignments (or potentially deprecation) for IEEE1621 Power State Set will be administered by IANA through Expert Review [RFC5226], i.e., review by one of a group of experts designated by an IETF Area Director. The group of experts MUST check the requested state for completeness and accuracy of the description.

12.2.2. IANA Registration of the DMTF Power State Set

This document specifies a set of values for the DMTF Power State Set. The list of these values with their identifiers is given in Section 5.2.1. The Internet Assigned Numbers Authority (IANA) has created a new registry for DMTF Power State Set identifiers and filled it with the initial list in Section 5.2.1.

New assignments (or potentially deprecation) for DMTF Power State Set will be administered by IANA through Expert Review [RFC5226], i.e., review by one of a group of experts designated by an IETF Area Director. The group of experts MUST check the conformance with the DMTF standard [DMTF], on the top of

12.2.3. IANA Registration of the EMAN Power State Set

This document specifies a set of values for the EMAN Power State Set. The list of these values with their identifiers is given in Section 5.2.1. The Internet Assigned Numbers Authority (IANA) has created a new registry for EMAN Power State Set identifiers and filled it with the initial list in Section 5.2.1.

New assignments (or potentially deprecation) for EMAN Power State Set New assignments in Power State Set require a Standards Action , i.e., they are to be made via Standards Track RFCs approved by the IESG.

12. Contributors

This document results from the merger of two initial proposals. The following persons made significant contributions either in one of the initial proposals or in this document.

John Parello

Rolf Winter

Dominique Dudkowski

13. Acknowledgment

The authors would like to thank Shamita Pisal for her prototype of this MIB module, and her valuable feedback. The authors would like to Michael Brown for improving the text dramatically.

14. Open Issues

OPEN ISSUE : double-check all the IEC references in the draft.

OPEN ISSUE: Description clause of pmPowerIndex Do we need this text Juergen Quittek to comment:

"The identity provisioning method that has been chosen can be retrieved by reading the value of powerStateEnergyConsumerOid. In case of identities provided by the ENERGY-AWARE-MIB module, this OID points to an existing instance of pmPowerIndex, in case of the ENTITY MIB, the object points to a valid instance of entPhysicalIndex, and in a similar way, it points to a value of another MIB module if this is used for identifying entities. If no other MIB module has been chosen for providing entity identities, then the value of powerStateEnergyConsumerOid MUST be 0.0 (zeroDotZero).

OPEN ISSUE : Juergen Schoenwalder review comments email May 25, 2011

PowerStateSeries ::= TEXTUAL-CONVENTION

Why is this an OCTET STRING (SIZE(1)) and not simply an enumerated INTEGER? And if this is to be maintained by IANA, why not create a IANA-POWER-SERIES-TC MIB module so that one can simply fetch the latest version from IANA?

New assignments in Power State Series require a Standards Action [RFC5226], i.e., they are to be made via Standards Track RFCs approved by the IESG.

This raises the bar pretty high. If some future organization defines popular power states, do you think someone is going to go through the trouble of producing a standards-track specification for this?

I also do not see why all objects in the pmPowerEntry are necessarily indexed by power series - some appear to me to be rather a property of the monitor and not the power state series the monitor happens to support.

Since I started looking at the IANA considerations, I believe this text needs to be removed:

OPEN ISSUE : Michael Schroff email comments Feb 24, 2011

TimeStamps for Power measurements

AC Power, Voltage, current measurement terminology

3-phase WYE or Delta or hybrid of WYE and Delta

Circuit breakers in scope of EMAN

Response sent to mailing list requesting for more information
and Clarification June 29, 2011.

15. References

15.2. Normative References

- [RFC2119] S. Bradner, Key words for use in RFCs to Indicate Requirement Levels, BCP 14, RFC 2119, March 1997.
- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.
- [RFC2579] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Textual Conventions for SMIv2", STD 58, RFC 2579, April 1999.
- [RFC2580] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Conformance Statements for SMIv2", STD 58, RFC 2580, April 1999.
- [RFC3621] Berger, A., and D. Romascanu, "Power Ethernet MIB", RFC3621, December 2003.
- [RFC4133] Bierman, A. and K. McCloghrie, "Entity MIB (Version 3)", RFC 4133, August 2005.
- [LLDP-MED-MIB] ANSI/TIA-1057, "The LLDP Management Information Base extension module for TIA-TR41.4 media endpoint discovery information", July 2005.
- [EMAN-AWARE-MIB] J. Parello, and B. Claise, "draft-ietf-eman-energy-aware-mib-02 ", work in progress, July 2011.

15.3. Informative References

- [RFC1628] S. Bradner, "UPS Management Information Base", RFC 1628, May 1994

Internet-Draft <Energy Monitoring MIB> July 2011

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

[RFC3418] Presun, R., Case, J., McCloghrie, K., Rose, M, and S.
 Waldbusser, "Management Information Base (MIB) for the
 Simple Network Management Protocol (SNMP)", RFC3418,
 December 2002.

[RFC3433] Bierman, A., Romascanu, D., and K. Norseth, "Entity
 Sensor Management Information Base", RFC 3433, December
 2002.

[RFC4268] Chisholm, S. and D. Perkins, "Entity State MIB", RFC
 4268, November 2005.

[RFC5226] Narten, T. Alverstrand, H., A. and K. McCloghrie,
 "Guidelines for Writing an IANA Considerations Section
 in RFCs ", BCP 26, RFC 5226, May 2008.

[EMAN-REQ] Quittek, J., Winter, R., Dietz, T., Claise, B., and
 M. Chandramouli, " Requirements for Energy Management
 ", draft-ietf-eman-requirements-03 (work in
 progress), July 2011. .

[EMAN-FRAMEWORK] Claise, B., Parello, J., Schoening, B., and J.
 Quittek, "Energy Management Framework", draft-ietf-
 eman-framework-02 , July 2011.

[EMAN-MONITORING-MIB] M. Chandramouli, Schoening, B., Dietz, T.,
 Quittek, J. and B. Claise "Energy and Power Monitoring
 MIB ", draft-claise-energy-monitoring-mib-09, July
 2011.

[EMAN-AS] Tychon, E., Laherty, M., and B. Schoening, "Energy
 Management (EMAN) Applicability Statement", draft-
 tychon-eman-applicability-statement-02, work in
 progress, June 2011.

[ACPI] "Advanced Configuration and Power Interface
 Specification", <http://www.acpi.info/DOWNLOADS/ACPIspec30b.pdf>

[DMTF] "Power State Management Profile DMTF DSP1027 Version
 2.0" December 2009

- [IEEE1621] "Standard for User Interface Elements in Power Control of Electronic Devices Employed in Office/Consumer Environments", IEEE 1621, December 2004.
- [IEC.61850-7-4] International Electrotechnical Commission, "Communication networks and systems for power utility automation Part 7-4: Basic communication structure Compatible logical node classes and data object classes", 2010.
- [IEC.62053-21] International Electrotechnical Commission, "Electricity metering equipment (a.c.) Particular requirements Part 22: Static meters for active energy (classes 1 and 2)", 2003.
- [IEC.62053-22] International Electrotechnical Commission, "Electricity metering equipment (a.c.) Particular requirements Part 22: Static meters for active energy (classes 0,2 S and 0,5 S)", 2003.

Authors' Addresses

Mouli Chandramouli
Cisco Systems, Inc.
Sarjapur Outer Ring Road
Bangalore,
IN

Phone: +91 80 4426 3947
Email: moulchan@cisco.com

Brad Schoening
44 Rivers Edge Drive
Little Silver, NJ 07739
US
Email: brad@bradschoening.com

Juergen Quittek
NEC Europe Ltd.

Internet-Draft <Energy Monitoring MIB>
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

July 2011

Phone: +49 6221 4342-115
Email: quittek@neclab.eu

Thomas Dietz
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-128
Email: Thomas.Dietz@neclab.eu

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

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

Network Working Group
Internet-Draft
Intended Status: Informational
Expires: April 20, 2011

B. Claise
J. Parello
B. Schoening
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/lid-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.

<Claise, et. Al>

Expires April 20 2011

[Page 1]

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.

1. Introduction.....	4
2. Use Cases & Requirements.....	5
3. Terminology.....	5
4. Energy Management Reference Model.....	7
5. Architecture High Level Concepts and Scope.....	9
5.1. Power Monitor Information.....	11
5.2. Power Monitor Meter Domain.....	11
5.3. Power Monitor Parent and Child.....	11
5.4. Power Monitor Context.....	12
5.5. Power Monitor Levels.....	13
5.6. Power Monitor Usage Measurement.....	16
5.7. Optional Power Usage Quality.....	17
5.8. Optional Energy Measurement.....	18
5.9. Optional Battery Information.....	18
6. Power Monitor Children Discovery.....	18
7. Configuration.....	19
8. Fault Management.....	20
9. IPFIX.....	20
10. Relationship with Other Standards Development Organizations.....	21
10.1. Information Modeling.....	21
10.2. Power Levels.....	21
11. Implementation Scenarios.....	22
Scenario 1: Switch with PoE endpoints.....	22
Scenario 2: Switch with PoE endpoints with further connected device(s).....	22
Scenario 3: A switch with Wireless Access Points.....	22
Scenario 4: Network connected facilities gateway.....	23
Scenario 5: Data center network.....	23
Scenario 6: Building gateway device.....	23
Scenario 7: Power consumption of UPS.....	23
Scenario 8: Power consumption of battery-based devices...	24
12. Security Considerations.....	24
12.1. Security Considerations for SNMP.....	24
12.2. Security Considerations for IPFIX.....	25
13. IANA Considerations.....	25
14. Acknowledgments.....	25
15. References.....	25
Normative References.....	25
Informative References.....	26

- 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: powerExample: 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

Internet-Draft <Power Management Archictechure> Octobre 2010
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

Internet-Draft <Power Management Archictechure> Octobre 2010
 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.

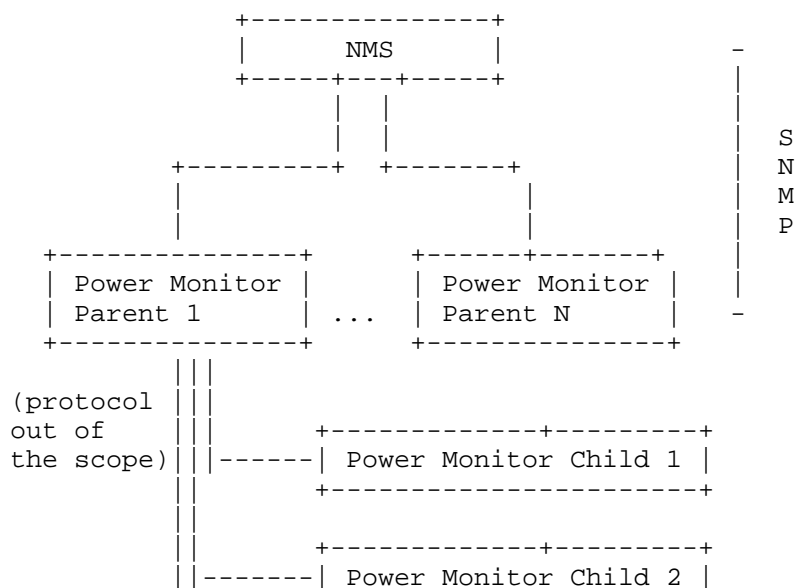
Power Level


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.

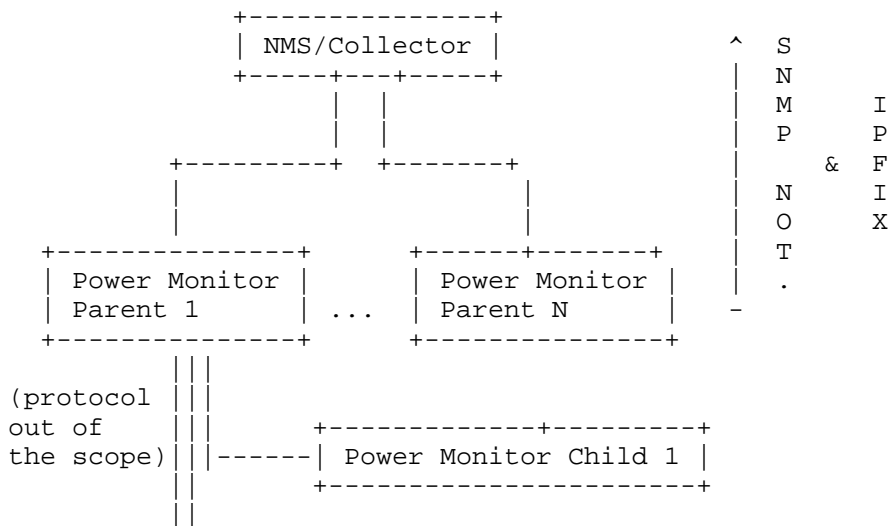
4. Energy Management Reference Model





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 be 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.






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

Internet-Draft <Power Management Architecture> Octobre 2010
- 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
-------	-----------------------------	---------------------

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

Internet-Draft <Power Management Archiectecure> Octobre 2010
 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 $\text{BaseValue} * 10$ 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

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].

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].

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.

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).

Internet-Draft <Power Management Architecture> Octobre 2010
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",

Internet-Draft <Power Management Archictechure> Octobre 2010
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

Internet-Draft <Power Management Archiectecure> Octobre 2010
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

Operations and Management Area Working Group
Internet Draft
Intended status: Informational
Expires: September 2014

T. Mizrahi
Marvell
N. Sprecher
Nokia Solutions and Networks
E. Bellagamba
Ericsson
Y. Weingarten

March 28, 2014

An Overview of
Operations, Administration, and Maintenance (OAM) Tools
draft-ietf-opsawg-oam-overview-16.txt

Abstract

Operations, Administration, and Maintenance (OAM) is a general term that refers to a toolset for fault detection and isolation, and for performance measurement. Over the years various OAM tools have been defined for various layers in the protocol stack.

This document summarizes some of the OAM tools defined in the IETF in the context of IP unicast, MPLS, MPLS Transport Profile (MPLS-TP), pseudowires, and TRILL. This document focuses on tools for detecting and isolating failures in networks and for performance monitoring. Control and management aspects of OAM are outside the scope of this document. Network repair functions such as Fast Reroute (FRR) and protection switching, which are often triggered by OAM protocols, are also out of the scope of this document.

The target audience of this document includes network equipment vendors, network operators and standards development organizations, and can be used as an index to some of the main OAM tools defined in the IETF. This document provides a brief description of each of the OAM tools in the IETF. At the end of the document a list of the OAM toolsets and a list of the OAM functions are presented as a summary.

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/lid-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 28, 2014.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	4
1.1. Background	4
1.2. Target Audience.....	5
1.3. OAM-related Work in the IETF	6
1.4. Focusing on the Data Plane	7
2. Terminology	7
2.1. Abbreviations	7
2.2. Terminology used in OAM Standards	9
2.2.1. General Terms	9
2.2.2. Operations, Administration and Maintenance	9
2.2.3. Functions, Tools and Protocols	10
2.2.4. Data Plane, Control Plane and Management Plane	11
2.2.5. The Players	12
2.2.6. Proactive and On-demand Activation	12
2.2.7. Connectivity Verification and Continuity Checks ...	13
2.2.8. Connection Oriented vs. Connectionless Communication	14
2.2.9. Point-to-point vs. Point-to-multipoint Services ...	14

2.2.10. Failures	15
3. OAM Functions	16
4. OAM Tools in the IETF - a Detailed Description	16
4.1. IP Ping	17
4.2. IP Traceroute	17
4.3. Bidirectional Forwarding Detection (BFD)	18
4.3.1. Overview	18
4.3.2. Terminology	19
4.3.3. BFD Control	19
4.3.4. BFD Echo	19
4.4. MPLS OAM	20
4.4.1. LSP Ping	20
4.4.2. BFD for MPLS	21
4.4.3. OAM for Virtual Private Networks (VPN) over MPLS ..	21
4.5. MPLS-TP OAM	21
4.5.1. Overview	21
4.5.2. Terminology	22
4.5.3. Generic Associated Channel	24
4.5.4. MPLS-TP OAM Toolset	24
4.5.4.1. Continuity Check and Connectivity Verification	25
4.5.4.2. Route Tracing	25
4.5.4.3. Lock Instruct	25
4.5.4.4. Lock Reporting	25
4.5.4.5. Alarm Reporting	26
4.5.4.6. Remote Defect Indication	26
4.5.4.7. Client Failure Indication	26
4.5.4.8. Performance Monitoring	26
4.5.4.8.1. Packet Loss Measurement (LM)	26
4.5.4.8.2. Packet Delay Measurement (DM)	27
4.6. Pseudowire OAM	27
4.6.1. Pseudowire OAM using Virtual Circuit Connectivity	
Verification (VCCV)	27
4.6.2. Pseudowire OAM using G-ACh	29
4.6.3. Attachment Circuit - Pseudowire Mapping	29
4.7. OWAMP and TWAMP.....	29
4.7.1. Overview	29
4.7.2. Control and Test Protocols	30
4.7.3. OWAMP	31
4.7.4. TWAMP	31
4.8. TRILL	32
5. Summary	32
5.1. Summary of OAM Tools	32
5.2. Summary of OAM Functions	35
5.3. Guidance to Network Equipment Vendors	36
6. Security Considerations	36
7. IANA Considerations	37
8. Acknowledgments	37

9. References	37
9.1. Normative References	37
9.2. Informative References	37
Appendix A. List of OAM Documents	43
A.1. List of IETF OAM Documents	43
A.2. List of Selected Non-IETF OAM Documents	48

1. Introduction

OAM is a general term that refers to a toolset for detecting, isolating and reporting failures and for monitoring the network performance.

There are several different interpretations to the "OAM" acronym. This document refers to Operations, Administration and Maintenance, as recommended in Section 3 of [OAM-Def].

This document summarizes some of the OAM tools defined in the IETF in the context of IP unicast, MPLS, MPLS Transport Profile (MPLS-TP), pseudowires, and TRILL.

This document focuses on tools for detecting and isolating failures and for performance monitoring. Hence, this document focuses on the tools used for monitoring and measuring the data plane; control and management aspects of OAM are outside the scope of this document. Network repair functions such as Fast Reroute (FRR) and protection switching, which are often triggered by OAM protocols, are also out of the scope of this document.

1.1. Background

OAM was originally used in traditional communication technologies such as E1 and T1, evolving into PDH and then later in SONET/SDH. ATM was probably the first technology to include inherent OAM support from day one, while in other technologies OAM was typically defined in an ad hoc manner after the technology was already defined and deployed. Packet-based networks were traditionally considered unreliable and best-effort. As packet-based networks evolved, they have become the common transport for both data and telephony, replacing traditional transport protocols. Consequently, packet-based networks were expected to provide a similar "carrier grade" experience, and specifically to support more advanced OAM functions, beyond ICMP and router hellos, that were traditionally used for fault detection.

As typical networks have a multi-layer architecture, the set of OAM protocols similarly take a multi-layer structure; each layer has its

own OAM protocols. Moreover, OAM can be used at different levels of hierarchy in the network to form a multi-layer OAM solution, as shown in the example in Figure 1.

Figure 1 illustrates a network in which IP traffic between two customer edges is transported over an MPLS provider network. MPLS OAM is used at the provider-level for monitoring the connection between the two provider edges, while IP OAM is used at the customer-level for monitoring the end-to-end connection between the two customer edges.




Figure 1 Example: Multi-layer OAM

1.2. Target Audience

The target audience of this document includes:

- o Standards development organizations - both IETF working groups and non-IETF organizations can benefit from this document when designing new OAM protocols, or when looking to reuse existing OAM tools for new technologies.
- o Network equipment vendors and network operators - can use this document as an index to some of the common IETF OAM tools.

It should be noted that some background in OAM is necessary in order to understand and benefit from this document. Specifically, the reader is assumed to be familiar with the term OAM [OAM-Def], the motivation for using OAM, and the distinction between OAM and network management [OAM-Mng].

1.3. OAM-related Work in the IETF

This memo provides an overview of the different sets of OAM tools defined by the IETF. The set of OAM tools described in this memo are applicable to IP unicast, MPLS, pseudowires, MPLS Transport Profile (MPLS-TP), and TRILL. While OAM tools that are applicable to other technologies exist, they are beyond the scope of this memo.

This document focuses on IETF documents that have been published as RFCs, while other ongoing OAM-related work is outside the scope.

The IETF has defined OAM protocols and tools in several different contexts. We roughly categorize these efforts into a few sets of OAM-related RFCs, listed in Table 1. Each set defines a logically-coupled set of RFCs, although the sets are in some cases intertwined by common tools and protocols.

The discussion in this document is ordered according to these sets (the acronyms and abbreviations are listed in Section 2.1.).

Toolset	Transport Technology
IP Ping	IPv4/IPv6
IP Traceroute	IPv4/IPv6
BFD	generic
MPLS OAM	MPLS
MPLS-TP OAM	MPLS-TP
Pseudowire OAM	Pseudowires
OWAMP and TWAMP	IPv4/IPv6
TRILL OAM	TRILL

Table 1 OAM Toolset Packages in the IETF Documents

This document focuses on OAM tools that have been developed in the IETF. A short summary of some of the significant OAM standards that have been developed in other standard organizations is presented in Appendix A.2.

1.4. Focusing on the Data Plane

OAM tools may, and quite often do, work in conjunction with a control plane and/or management plane. OAM provides instrumentation tools for measuring and monitoring the data plane. OAM tools often use control plane functions, e.g., to initialize OAM sessions and to exchange various parameters. The OAM tools communicate with the management plane to raise alarms, and often OAM tools may be activated by the management (as well as by the control plane), e.g., to locate and localize problems.

The considerations of the control plane maintenance tools and the functionality of the management plane are out of scope for this document, which concentrates on presenting the data plane tools that are used for OAM. Network repair functions such as Fast Reroute (FRR) and protection switching, which are often triggered by OAM protocols, are also out of the scope of this document.

Since OAM protocols are used for monitoring the data plane, it is imperative for OAM tools to be capable of testing the actual data plane with as much accuracy as possible. Thus, it is important to enforce fate-sharing between OAM traffic that monitors the data plane and the data plane traffic it monitors.

2. Terminology

2.1. Abbreviations

ACH	Associated Channel Header
AIS	Alarm Indication Signal
ATM	Asynchronous Transfer Mode
BFD	Bidirectional Forwarding Detection
CC	Continuity Check
CV	Connectivity Verification
DM	Delay Measurement

ECMP	Equal Cost Multiple Paths
FEC	Forwarding Equivalence Class
FRR	Fast Reroute
G-ACh	Generic Associated Channel
GAL	Generic Associated Label
ICMP	Internet Control Message Protocol
L2TP	Layer Two Tunneling Protocol
L2VPN	Layer Two Virtual Private Network
L3VPN	Layer Three Virtual Private Network
LCCE	L2TP Control Connection Endpoint
LDP	Label Distribution Protocol
LER	Label Edge Router
LM	Loss Measurement
LSP	Label Switched Path
LSR	Label Switched Router
ME	Maintenance Entity
MEG	Maintenance Entity Group
MEP	MEG End Point
MIP	MEG Intermediate Point
MP	Maintenance Point
MPLS	Multiprotocol Label Switching
MPLS-TP	MPLS Transport Profile
MTU	Maximum Transmission Unit
OAM	Operations, Administration, and Maintenance

OWAMP	One-way Active Measurement Protocol
PDH	Plesiochronous Digital Hierarchy
PE	Provider Edge
PSN	Public Switched Network
PW	Pseudowire
PWE3	Pseudowire Emulation Edge-to-Edge
RBridge	Routing Bridge
RDI	Remote Defect Indication
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Networking
TRILL	Transparent Interconnection of Lots of Links
TTL	Time To Live
TWAMP	Two-way Active Measurement Protocol
VCCV	Virtual Circuit Connectivity Verification
VPN	Virtual Private Network

2.2. Terminology used in OAM Standards

2.2.1. General Terms

A wide variety of terms is used in various OAM standards. This section presents a comparison of the terms used in various OAM standards, without fully quoting the definition of each term.

An interesting overview of the term OAM and its derivatives is presented in [OAM-Def]. A thesaurus of terminology for MPLS-TP terms is presented in [TP-Term], and provides a good summary of some of the OAM related terminology.

2.2.2. Operations, Administration and Maintenance

The following definition of OAM is quoted from [OAM-Def]:

The components of the "OAM" acronym (and provisioning) are defined as follows:

- o Operations - Operation activities are undertaken to keep the network (and the services that the network provides) up and running. It includes monitoring the network and finding problems. Ideally these problems should be found before users are affected.
- o Administration - Administration activities involve keeping track of resources in the network and how they are used. It includes all the bookkeeping that is necessary to track networking resources and the network under control.
- o Maintenance - Maintenance activities are focused on facilitating repairs and upgrades -- for example, when equipment must be replaced, when a router needs a patch for an operating system image, or when a new switch is added to a network. Maintenance also involves corrective and preventive measures to make the managed network run more effectively, e.g., adjusting device configuration and parameters.

2.2.3. Functions, Tools and Protocols

OAM Function

An OAM function is an instrumentation measurement type or diagnostic.

OAM functions are the atomic building blocks of OAM, where each function defines an OAM capability.

Typical examples of OAM functions are presented in Section 3.

OAM Protocol

A protocol used for implementing one or more OAM functions.

The OWAMP-Test [OWAMP] is an example of an OAM protocol.

OAM Tool

An OAM tool is a specific means of applying one or more OAM functions.

In some cases an OAM protocol *is* an OAM tool, e.g., OWAMP-Test. In other cases an OAM tool uses a set of protocols that are not strictly OAM-related; for example, Traceroute (Section 4.2.) can be

implemented using UDP and ICMP messages, without using an OAM protocol per se.

2.2.4. Data Plane, Control Plane and Management Plane

Data Plane

The data plane is the set of functions used to transfer data in the stratum or layer under consideration [ITU-Terms].

The Data Plane is also known as the Forwarding Plane or the User Plane.

Control Plane

The control plane is the set of protocols and mechanisms that enable routers to efficiently learn how to forward packets towards their final destination (based on [Comp]).

Management Plane

The term Management Plane, as described in [Mng], is used to describe the exchange of management messages through management protocols (often transported by IP and by IP transport protocols) between management applications and the managed entities such as network nodes.

Data Plane vs. Control Plane vs. Management Plane

The distinction between the planes is at times a bit vague. For example, the definition of "Control Plane" above may imply that OAM tools such as ping, BFD and others are in fact in the control plane.

This document focuses on tools used for monitoring the data plane. While these tools could arguably be considered to be in the control plane, these tools monitor the data plane, and hence it is imperative to have fate-sharing between OAM traffic that monitors the data plane and the data plane traffic it monitors.

Another potentially vague distinction is between the management plane and control plane. The management plane should be seen as separate from, but possibly overlapping with, the control plane (based on [Mng]).

2.2.5. The Players

An OAM tool is used between two (or more) peers. Various terms are used in IETF documents to refer to the players that take part in OAM. Table 2 summarizes the terms used in each of the toolsets discussed in this document.

Toolset	Terms
Ping / Traceroute ([ICMPv4], [ICMPv6], [TCPIP-Tools])	-Host -Node -Interface -Gateway
BFD [BFD]	System
MPLS OAM [MPLS-OAM-FW]	LSR
MPLS-TP OAM [TP-OAM-FW]	-End Point - MEP -Intermediate Point - MIP
Pseudowire OAM [VCCV]	-PE -LCCE
OWAMP and TWAMP ([OWAMP], [TWAMP])	-Host -End system
TRILL OAM [TRILL-OAM]	-RBridge

Table 2 Maintenance Point Terminology

2.2.6. Proactive and On-demand Activation

The different OAM tools may be used in one of two basic types of activation:

Proactive

Proactive activation - indicates that the tool is activated on a continual basis, where messages are sent periodically, and errors are detected when a certain number of expected messages are not received.

On-demand

On-demand activation - indicates that the tool is activated "manually" to detect a specific anomaly.

2.2.7. Connectivity Verification and Continuity Checks

Two distinct classes of failure management functions are used in OAM protocols, connectivity verification and continuity checks. The distinction between these terms is defined in [MPLS-TP-OAM], and is used similarly in this document.

Continuity Check

Continuity checks are used to verify that a destination is reachable, and are typically sent proactively, though they can be invoked on-demand as well.

Connectivity Verification

A connectivity verification function allows Alice to check whether she is connected to Bob or not. It is noted that while the CV function is performed in the data plane, the "expected path" is predetermined either in the control plane or in the management plane. A connectivity verification (CV) protocol typically uses a CV message, followed by a CV reply that is sent back to the originator. A CV function can be applied proactively or on-demand.

Connectivity verification tools often perform path verification as well, allowing Alice to verify that messages from Bob are received through the correct path, thereby verifying not only that the two MPs are connected, but also that they are connected through the expected path, allowing detection of unexpected topology changes.

Connectivity verification functions can also be used for checking the MTU of the path between the two peers.

Connectivity verification and continuity checks are considered complementary mechanisms, and are often used in conjunction with each other.

2.2.8. Connection Oriented vs. Connectionless Communication

Connection Oriented

In Connection Oriented technologies an end-to-end connection is established (by a control protocol or provisioned by a management system) prior to the transmission of data.

Typically a connection identifier is used to identify the connection. In connection oriented technologies it is often the case (although not always) that all packets belonging to a specific connection use the same route through the network.

Connectionless

In Connectionless technologies data is typically sent between end points without prior arrangement. Packets are routed independently based on their destination address, and hence different packets may be routed in a different way across the network.

Discussion

The OAM tools described in this document include tools that support connection oriented technologies, as well as tools for connectionless technologies.

In connection oriented technologies OAM is used to monitor a *specific* connection; OAM packets are forwarded through the same route as the data traffic and receive the same treatment. In connectionless technologies, OAM is used between a source and destination pair without defining a specific connection. Moreover, in some cases the route of OAM packets may differ from the one of the data traffic. For example, the connectionless IP Ping (Section 4.1.) tests the reachability from a source to a given destination, while the connection oriented LSP Ping (Section 4.4.) is used for monitoring a specific LSP (connection), and provides the capability to monitor all the available paths used by an LSP.

It should be noted that in some cases connectionless protocols are monitored by connection oriented OAM protocols. For example, while IP is a connectionless protocol, it can be monitored by BFD (Section 4.3.), which is connection oriented.

2.2.9. Point-to-point vs. Point-to-multipoint Services

Point-to-point (P2P)

A P2P service delivers data from a single source to a single destination.

Point-to-multipoint (P2MP)

A P2MP service delivers data from a single source to a one or more destinations (based on [Signal]).

An MP2MP service is a service that delivers data from more than one source to one or more receivers (based on [Signal]).

Note: the two definitions for P2MP and MP2MP are quoted from [Signal]. Although [Signal] describes a specific case of P2MP and MP2MP which is MPLS-specific, these two definitions also apply to non-MPLS cases.

Discussion

The OAM tools described in this document include tools for P2P services, as well as tools for P2MP services.

The distinction between P2P services and P2MP services affects the corresponding OAM tools. A P2P service is typically simpler to monitor, as it consists of a single pair of end points. P2MP and MP2MP services present several challenges. For example, in a P2MP service, the OAM mechanism not only verifies that each of the destinations is reachable from the source, but also verifies that the P2MP distribution tree is intact and loop-free.

2.2.10. Failures

The terms Failure, Fault, and Defect are used interchangeably in the standards, referring to a malfunction that can be detected by a connectivity or a continuity check. In some standards, such as 802.1ag [IEEE802.1Q], there is no distinction between these terms, while in other standards each of these terms refers to a different type of malfunction.

The terminology used in IETF MPLS-TP OAM is based on the ITU-T terminology, which distinguishes between these three terms in [ITU-T-G.806];

Fault

The term Fault refers to an inability to perform a required action, e.g., an unsuccessful attempt to deliver a packet.

Defect

The term Defect refers to an interruption in the normal operation, such as a consecutive period of time where no packets are delivered successfully.

Failure

The term Failure refers to the termination of the required function. While a Defect typically refers to a limited period of time, a failure refers to a long period of time.

3. OAM Functions

This subsection provides a brief summary of the common OAM functions used in OAM-related standards. These functions are used as building blocks in the OAM standards described in this document.

- o Connectivity Verification (CV), Path Verification and Continuity Checks (CC):
As defined in Section 2.2.7.
- o Path Discovery / Fault Localization:
This function can be used to trace the route to a destination, i.e., to identify the nodes along the route to the destination. When more than one route is available to a specific destination, this function traces one of the available routes. When a failure occurs, this function attempts to detect the location of the failure.
Note that the term route tracing (or Traceroute) that is used in the context of IP and MPLS, is sometimes referred to as path tracing in the context of other protocols, such as TRILL.
- o Performance Monitoring:
Typically refers to:
 - o Loss Measurement (LM) - monitors the packet loss rate.
 - o Delay Measurement (DM) - monitors the delay and delay variation (jitter).

4. OAM Tools in the IETF - a Detailed Description

This section presents a detailed description of the sets of OAM-related tools in each of the toolsets in Table 1.

4.1. IP Ping

Ping is a common network diagnosis application for IP networks that uses ICMP. According to [NetTerms], 'Ping' is an abbreviation for Packet internet groper, although the term has been so commonly used that it stands on its own. As defined in [NetTerms], it is a program used to test reachability of destinations by sending them an ICMP echo request and waiting for a reply.

The ICMP Echo request/reply exchange in Ping is used as a continuity check function for the Internet Protocol. The originator transmits an ICMP Echo request packet, and the receiver replies with an Echo reply. ICMP ping is defined in two variants, [ICMPv4] is used for IPv4, and [ICMPv6] is used for IPv6.

Ping can be invoked either to a unicast destination or to a multicast destination. In the latter case, all members of the multicast group send an Echo reply back to the originator.

Ping implementations typically use ICMP messages. UDP Ping is a variant that uses UDP messages instead of ICMP echo messages.

Ping is a single-ended continuity check, i.e., it allows the *initiator* of the Echo request to test the reachability. If it is desirable for both ends to test the reachability, both ends have to invoke Ping independently.

Note that since ICMP filtering is deployed in some routers and firewalls, the usefulness of Ping is sometimes limited in the wider internet. This limitation is equally relevant to Traceroute.

4.2. IP Traceroute

Traceroute ([TCPIP-Tools], [NetTools]) is an application that allows users to discover a path between an IP source and an IP destination.

The most common way to implement Traceroute [TCPIP-Tools] is described as follows. Traceroute sends a sequence of UDP packets to UDP port 33434 at the destination. By default, Traceroute begins by sending three packets (the number of packets is configurable in most Traceroute implementations), each with an IP Time-To-Live (or Hop Limit in IPv6) value of one to the destination. These packets expire as soon as they reach the first router in the path. Consequently, that router sends three ICMP Time Exceeded Messages back to the Traceroute application. Traceroute now sends another three UDP packets, each with the TTL value of 2. These messages cause the second router to return ICMP messages. This process continues, with

ever increasing values for the TTL field, until the packets actually reach the destination. Because no application listens to port 33434 at the destination, the destination returns ICMP Destination Unreachable Messages indicating an unreachable port. This event indicates to the Traceroute application that it is finished. The Traceroute program displays the round-trip delay associated with each of the attempts.

While Traceroute is a tool that finds *a* path from A to B, it should be noted that traffic from A to B is often forwarded through Equal Cost Multiple Paths (ECMP). Paris Traceroute [PARIS] is an extension to Traceroute that attempts to discover all the available paths from A to B by scanning different values of header fields (such as UDP ports) in the probe packets.

It is noted that Traceroute is an application, and not a protocol. As such, it has various different implementations. One of the most common ones uses UDP probe packets, as described above. Other implementations exist that use other types of probe messages, such as ICMP or TCP.

Note that IP routing may be asymmetric. While Traceroute discovers a path between a source and destination, it does not reveal the reverse path.

A few ICMP extensions ([ICMP-MP], [ICMP-Int]) have been defined in the context of Traceroute. These documents define several extensions, including extensions to the ICMP Destination Unreachable message, that can be used by Traceroute applications.

Traceroute allows path discovery to *unicast* destination addresses. A similar tool [mtrace] was defined for multicast destination addresses, allowing to trace the route that a multicast IP packet takes from a source to a particular receiver.

4.3. Bidirectional Forwarding Detection (BFD)

4.3.1. Overview

While multiple OAM tools have been defined for various protocols in the protocol stack, Bidirectional Forwarding Detection [BFD], defined by the IETF BFD working group, is a generic OAM tool that can be deployed over various encapsulating protocols, and in various medium types. The IETF has defined variants of the protocol for IP ([BFD-IP], [BFD-Multi]), for MPLS LSPs [BFD-LSP], and for pseudowires [BFD-VCCV]. The usage of BFD in MPLS-TP is defined in [TP-CC-CV].

BFD includes two main OAM functions, using two types of BFD packets: BFD Control packets, and BFD Echo packets.

4.3.2. Terminology

BFD operates between **systems**. The BFD protocol is run between two or more systems after establishing a **session**.

4.3.3. BFD Control

BFD supports a bidirectional continuity check, using BFD control packets, that are exchanged within a BFD session. BFD sessions operate in one of two modes:

- o Asynchronous mode (i.e., proactive): in this mode BFD control packets are sent periodically. When the receiver detects that no BFD control packets have been received during a predetermined period of time, a failure is reported.
- o Demand mode: in this mode, BFD control packets are sent on-demand. Upon need, a system initiates a series of BFD control packets to check the continuity of the session. BFD control packets are sent independently in each direction.

Each of the end-points (referred to as systems) of the monitored path maintains its own session identification, called a Discriminator, both of which are included in the BFD Control Packets that are exchanged between the end-points. At the time of session establishment, the Discriminators are exchanged between the two-end points. In addition, the transmission (and reception) rate is negotiated between the two end-points, based on information included in the control packets. These transmission rates may be renegotiated during the session.

During normal operation of the session, i.e., when no failures have been detected, the BFD session is in the Up state. If no BFD Control packets are received during a period of time called the Detection Time, the session is declared to be Down. The detection time is a function of the pre-configured or negotiated transmission rate, and a parameter called Detect Mult. Detect Mult determines the number of missing BFD Control packets that cause the session to be declared as Down. This parameter is included in the BFD Control packet.

4.3.4. BFD Echo

A BFD echo packet is sent to a peer system, and is looped back to the originator. The echo function can be used proactively, or on-demand.

The BFD echo function has been defined in BFD for IPv4 and IPv6 ([BFD-IP]), but is not used in BFD for MPLS LSPs, PWs, or in BFD for MPLS-TP.

4.4. MPLS OAM

The IETF MPLS working group has defined OAM for MPLS LSPs. The requirements and framework of this effort are defined in [MPLS-OAM-FW] and [MPLS-OAM], respectively. The corresponding OAM tool defined, in this context, is LSP Ping [LSP-Ping]. OAM for P2MP services is defined in [MPLS-P2MP].

BFD for MPLS [BFD-LSP] is an alternative means for detecting data-plane failures, as described below.

4.4.1. LSP Ping

LSP Ping is modeled after the Ping/Traceroute paradigm and thus it may be used in one of two modes:

- o "Ping" mode: In this mode LSP Ping is used for end-to-end connectivity verification between two LERs.
- o "Traceroute" mode: This mode is used for hop-by-hop fault isolation.

LSP Ping is based on ICMP Ping operation (of data-plane connectivity verification) with additional functionality to verify data-plane vs. control-plane consistency for a Forwarding Equivalence Class (FEC) and also identify Maximum Transmission Unit (MTU) problems.

The Traceroute functionality may be used to isolate and localize MPLS faults, using the Time-to-live (TTL) indicator to incrementally identify the sub-path of the LSP that is successfully traversed before the faulty link or node.

The challenge in MPLS networks is that the traffic of a given LSP may be load balanced across Equal Cost Multiple paths (ECMP). LSP Ping monitors all the available paths of an LSP by monitoring its different Forwarding Equivalence Classes (FEC). Note that MPLS-TP does not use ECMP, and thus does not require OAM over multiple paths.

Another challenge is that an MPLS LSP does not necessarily have a return path; traffic that is sent back from the egress LSR to the ingress LSR is not necessarily sent over an MPLS LSP, but can be sent through a different route, such as an IP route. Thus, responding to an LSP Ping message is not necessarily as trivial as in IP Ping,

where the responder just swaps the source and destination IP addresses. Note that this challenge is not applicable to MPLS-TP, where a return path is always available.

It should be noted that LSP Ping supports unique identification of the LSP within an addressing domain. The identification is checked using the full FEC identification. LSP Ping is extensible to include additional information needed to support new functionality, by use of Type-Length-Value (TLV) constructs. The usage of TLVs is typically handled by the control plane, as it is not easy to implement in hardware.

LSP Ping supports both asynchronous, as well as, on-demand activation.

4.4.2. BFD for MPLS

BFD [BFD-LSP] can be used to detect MPLS LSP data plane failures.

A BFD session is established for each MPLS LSP that is being monitored. BFD Control packets must be sent along the same path as the monitored LSP. If the LSP is associated with multiple FECs, a BFD session is established for each FEC.

While LSP Ping can be used for detecting MPLS data plane failures and for verifying the MPLS LSP data plane against the control plane, BFD can only be used for the former. BFD can be used in conjunction with LSP Ping, as is the case in MPLS-TP (see Section 4.5.4.).

4.4.3. OAM for Virtual Private Networks (VPN) over MPLS

The IETF has defined two classes of VPNs, Layer 2 VPNs (L2VPN) and Layer 3 VPNs (L3VPN). [L2VPN-OAM] provides the requirements and framework for OAM in the context of Layer 2 Virtual Private Networks (L2VPN), and specifically it also defines the OAM layering of L2VPNs over MPLS. [L3VPN-OAM] provides a framework for the operation and management of Layer 3 Virtual Private Networks (L3VPNs).

4.5. MPLS-TP OAM

4.5.1. Overview

The MPLS working group has defined the OAM toolset that fulfills the requirements for MPLS-TP OAM. The full set of requirements for MPLS-TP OAM are defined in [MPLS-TP-OAM], and include both general requirements for the behavior of the OAM tools and a set of operations that should be supported by the OAM toolset. The set of

mechanisms required are further elaborated in [TP-OAM-FW], which describes the general architecture of the OAM system as well as giving overviews of the functionality of the OAM toolset.

Some of the basic requirements for the OAM toolset for MPLS-TP are:

- o MPLS-TP OAM must be able to support both an IP based and non-IP based environment. If the network is IP based, i.e., IP routing and forwarding are available, then the MPLS-TP OAM toolset should rely on the IP routing and forwarding capabilities. On the other hand, in environments where IP functionality is not available, the OAM tools must still be able to operate without dependence on IP forwarding and routing.
- o OAM packets and the user traffic are required to be congruent (i.e., OAM packets are transmitted in-band) and there is a need to differentiate OAM packets from ordinary user packets in the data plane. Inherent in this requirement is the principle that MPLS-TP OAM be independent of any existing control-plane, although it should not preclude use of the control-plane functionality. OAM packets are identified by the Generic Associated Label (GAL), which is a reserved MPLS label value (13).

4.5.2. Terminology

Maintenance Entity (ME)

The MPLS-TP OAM tools are designed to monitor and manage a Maintenance Entity (ME). An ME, as defined in [TP-OAM-FW], defines a relationship between two points of a transport path to which maintenance and monitoring operations apply.

The term Maintenance Entity (ME) is used in ITU-T Recommendations (e.g., [ITU-T-Y1731]), as well as in the MPLS-TP terminology ([TP-OAM-FW]).

Maintenance Entity Group (MEG)

The collection of one or more MEs that belongs to the same transport path and that are maintained and monitored as a group are known as a Maintenance Entity Group (based on [TP-OAM-FW]).

Maintenance Point (MP)

A Maintenance Point (MP) is a functional entity that is defined at a node in the network, and can initiate and/or react to OAM messages. This document focuses on the data-plane functionality of MPs, while

MPs interact with the control plane and with the management plane as well.

The term MP is used in IEEE 802.1ag, and was similarly adopted in MPLS-TP ([TP-OAM-FW]).

Maintenance End Point (MEP)

A Maintenance End Point (MEP) is one of the end points of an ME, and can initiate OAM messages and respond to them (based on [TP-OAM-FW]).

Maintenance Intermediate Point (MIP)

In between MEPs, there are zero or more intermediate points, called Maintenance Entity Group Intermediate Points (based on [TP-OAM-FW]).

A Maintenance Intermediate Point (MIP) is an intermediate point that does not generally initiate OAM frames (one exception to this is the use of AIS notifications), but is able to respond to OAM frames that are destined to it. A MIP in MPLS-TP identifies OAM packets destined to it by the expiration of the TTL field in the OAM packet. The term Maintenance Point is a general term for MEPs and MIPs.

Up and Down MEPs

The IEEE 802.1ag [IEEE802.1Q] defines a distinction between Up MEPs and Down MEPs. A MEP monitors traffic either in the direction facing the network, or in the direction facing the bridge. A Down MEP is a MEP that receives OAM packets from, and transmits them to the direction of the network. An Up MEP receives OAM packets from, and transmits them to the direction of the bridging entity. MPLS-TP ([TP-OAM-FW]) uses a similar distinction on the placement of the MEP - either at the ingress, egress, or forwarding function of the node (Down / Up MEPs). This placement is important for localization of a failure.

Note that the terms Up and Down MEPs are entirely unrelated to the conventional up/down terminology, where down means faulty, and up is nonfaulty.

The distinction between Up and Down MEPs was defined in [TP-OAM-FW], but has not been used in other MPLS-TP RFCs, as of the writing of this document.

4.5.3. Generic Associated Channel

In order to address the requirement for in-band transmission of MPLS-TP OAM traffic, MPLS-TP uses a Generic Associated Channel (G-ACh), defined in [G-ACh] for LSP-based OAM traffic. This mechanism is based on the same concepts as the PWE3 ACH [PW-ACH] and VCCV [VCCV] mechanisms. However, to address the needs of LSPs as differentiated from PW, the following concepts were defined for [G-ACh]:

- o An Associated Channel Header (ACH), that uses a format similar to the PW Control Word [PW-ACH], is a 4-byte header that is prepended to OAM packets.
- o A Generic Associated Label (GAL). The GAL is a reserved MPLS label value (13) that indicates that the packet is an ACH packet and the payload follows immediately after the label stack.

It should be noted that while the G-ACh was defined as part of the MPLS-TP definition effort, the G-ACh is a generic tool that can be used in MPLS in general, and not only in MPLS-TP.

4.5.4. MPLS-TP OAM Toolset

To address the functionality that is required of the OAM toolset, the MPLS WG conducted an analysis of the existing IETF and ITU-T OAM tools and their ability to fulfill the required functionality. The conclusions of this analysis are documented in [OAM-Analys]. MPLS-TP uses a mixture of OAM tools that are based on previous standards, and adapted to the requirements of [MPLS-TP-OAM]. Some of the main building blocks of this solution are based on:

- o Bidirectional Forwarding Detection ([BFD], [BFD-LSP]) for proactive continuity check and connectivity verification.
- o LSP Ping as defined in [LSP-Ping] for on-demand connectivity verification.
- o New protocol packets, using G-ACh, to address different functionality.
- o Performance measurement protocols that are based on the functionality that is described in [ITU-T-Y1731].

The following sub-sections describe the OAM tools defined for MPLS-TP as described in [TP-OAM-FW].

4.5.4.1. Continuity Check and Connectivity Verification

Continuity Check and Connectivity Verification are presented in Section 2.2.7. of this document. As presented there, these tools may be used either proactively or on-demand. When using these tools proactively, they are generally used in tandem.

For MPLS-TP there are two distinct tools, the proactive tool is defined in [TP-CC-CV] while the on-demand tool is defined in [OnDemand-CV]. In on-demand mode, this function should support monitoring between the MEPs and, in addition, between a MEP and MIP. [TP-OAM-FW] highlights, when performing Connectivity Verification, the need for the CC-V messages to include unique identification of the MEG that is being monitored and the MEP that originated the message.

The proactive tool [TP-CC-CV] is based on extensions to BFD (see Section 4.3.) with the additional limitation that the transmission and receiving rates are based on configuration by the operator. The on-demand tool [OnDemand-CV] is an adaptation of LSP Ping (see Section 4.4.) for the required behavior of MPLS-TP.

4.5.4.2. Route Tracing

[MPLS-TP-OAM] defines that there is a need for functionality that would allow a path end-point to identify the intermediate and end-points of the path. This function would be used in on-demand mode. Normally, this path will be used for bidirectional PW, LSP, and sections, however, unidirectional paths may be supported only if a return path exists. The tool for this is based on the LSP Ping (see Section 4.4.) functionality and is described in [OnDemand-CV].

4.5.4.3. Lock Instruct

The Lock Instruct function [Lock-Loop] is used to notify a transport path end-point of an administrative need to disable the transport path. This functionality will generally be used in conjunction with some intrusive OAM function, e.g., Performance measurement, Diagnostic testing, to minimize the side-effect on user data traffic.

4.5.4.4. Lock Reporting

Lock Reporting is a function used by an end-point of a path to report to its far-end end-point that a lock condition has been affected on the path.

4.5.4.5. Alarm Reporting

Alarm Reporting [TP-Fault] provides the means to suppress alarms following detection of defect conditions at the server sub-layer. Alarm reporting is used by an intermediate point of a path, that becomes aware of a fault on the path, to report to the end-points of the path. [TP-OAM-FW] states that this may occur as a result of a defect condition discovered at a server sub-layer. This generates an Alarm Indication Signal (AIS) that continues until the fault is cleared. The consequent action of this function is detailed in [TP-OAM-FW].

4.5.4.6. Remote Defect Indication

Remote Defect Indication (RDI) is used proactively by a path end-point to report to its peer end-point that a defect is detected on a bidirectional connection between them. [MPLS-TP-OAM] points out that this function may be applied to a unidirectional LSP only if a return path exists. [TP-OAM-FW] points out that this function is associated with the proactive CC-V function.

4.5.4.7. Client Failure Indication

Client Failure Indication (CFI) is defined in [MPLS-TP-OAM] to allow the propagation information from one edge of the network to the other. The information concerns a defect to a client, in the case that the client does not support alarm notification.

4.5.4.8. Performance Monitoring

The definition of MPLS performance monitoring was motivated by the MPLS-TP requirements [MPLS-TP-OAM], but was defined generically for MPLS in [MPLS-LM-DM]. An additional document [TP-LM-DM] defines a performance monitoring profile for MPLS-TP.

4.5.4.8.1. Packet Loss Measurement (LM)

Packet Loss Measurement is a function used to verify the quality of the service. Packet loss, as defined in [IPPM-1LM] and [MPLS-TP-OAM], indicates the ratio of the number of user packets lost to the total number of user packets sent during a defined time interval.

There are two possible ways of determining this measurement:

- o Using OAM packets, it is possible to compute the statistics based on a series of OAM packets. This, however, has the disadvantage of being artificial, and may not be representative since part of the packet loss may be dependent upon packet sizes and upon the implementation of the MEPs that take part in the protocol.
- o Sending delimiting messages for the start and end of a measurement period during which the source and sink of the path count the packets transmitted and received. After the end delimiter, the ratio would be calculated by the path OAM entity.

4.5.4.8.2. Packet Delay Measurement (DM)

Packet Delay Measurement is a function that is used to measure one-way or two-way delay of a packet transmission between a pair of the end-points of a path (PW, LSP, or Section). Where:

- o One-way packet delay, as defined in [IPPM-1DM], is the time elapsed from the start of transmission of the first bit of the packet by a source node until the reception of the last bit of that packet by the destination node. Note that one-way delay measurement requires the clocks of the two end-points to be synchronized.
- o Two-way packet delay, as defined in [IPPM-2DM], is the time elapsed from the start of transmission of the first bit of the packet by a source node until the reception of the last bit of the loop-backed packet by the same source node, when the loopback is performed at the packet's destination node. Note that due to possible path asymmetry, the one-way packet delay from one end-point to another is not necessarily equal to half of the two-way packet delay.
As opposed to one-way delay measurement, two-way delay measurement does not require the two end-points to be synchronized.

For each of these two metrics, the DM function allows the MEP to measure the delay, as well as the delay variation. Delay measurement is performed by exchanging timestamped OAM packets between the participating MEPs.

4.6. Pseudowire OAM

4.6.1. Pseudowire OAM using Virtual Circuit Connectivity Verification (VCCV)

VCCV, as defined in [VCCV], provides a means for end-to-end fault detection and diagnostics tools to be used for PWs (regardless of the

underlying tunneling technology). The VCCV switching function provides a control channel associated with each PW. [VCCV] defines three Control Channel (CC) types, i.e., three possible methods for transmitting and identifying OAM messages:

- o CC Type 1: In-band VCCV, as described in [VCCV], is also referred to as "PWE3 Control Word with 0001b as first nibble". It uses the PW Associated Channel Header [PW-ACH].
- o CC Type 2: Out-of-band VCCV [VCCV], is also referred to as "MPLS Router Alert Label". In this case the control channel is created by using the MPLS router alert label [MPLS-ENCAPS] immediately above the PW label.
- o CC Type 3: TTL expiry VCCV [VCCV], is also referred to as "MPLS PW Label with TTL == 1", i.e., the control channel is identified when the value of the TTL field in the PW label is set to 1.

VCCV currently supports the following OAM tools: ICMP Ping, LSP Ping, and BFD. ICMP and LSP Ping are IP encapsulated before being sent over the PW ACH. BFD for VCCV [BFD-VCCV] supports two modes of encapsulation - either IP/UDP encapsulated (with IP/UDP header) or PW-ACH encapsulated (with no IP/UDP header) and provides support to signal the AC status. The use of the VCCV control channel provides the context, based on the MPLS-PW label, required to bind and bootstrap the BFD session to a particular pseudo wire (FEC), eliminating the need to exchange Discriminator values.

VCCV consists of two components: (1) signaled component to communicate VCCV capabilities as part of VC label, and (2) switching component to cause the PW payload to be treated as a control packet.

VCCV is not directly dependent upon the presence of a control plane. The VCCV capability advertisement may be performed as part of the PW signaling when LDP is used. In case of manual configuration of the PW, it is the responsibility of the operator to set consistent options at both ends. The manual option was created specifically to handle MPLS-TP use cases where no control plane was a requirement. However, new use cases such as pure mobile backhaul find this functionality useful too.

The PWE3 working group has conducted an implementation survey of VCCV [VCCV-SURVEY], which analyzes which VCCV mechanisms are used in practice.

4.6.2. Pseudowire OAM using G-ACh

As mentioned above, VCCV enables OAM for PWs by using a control channel for OAM packets. When PWs are used in MPLS-TP networks, rather than the control channels defined in VCCV, the G-ACh can be used as an alternative control channel. The usage of the G-ACh for PWs is defined in [PW-G-ACh].

4.6.3. Attachment Circuit - Pseudowire Mapping

The PWE3 working group has defined a mapping and notification of defect states between a pseudowire (PW) and the Attachment Circuits (ACs) of the end-to-end emulated service. This mapping is of key importance to the end-to-end functionality. Specifically, the mapping is provided by [PW-MAP], by [L2TP-EC] for L2TPv3 pseudowires, and Section 5.3 of [ATM-L2] for ATM.

[L2VPN-OAM] provides the requirements and framework for OAM in the context of Layer 2 Virtual Private Networks (L2VPN), and specifically it also defines the OAM layering of L2VPNs over pseudowires.

The mapping defined in [Eth-Int] allows an end-to-end emulated Ethernet service over pseudowires.

4.7. OWAMP and TWAMP

4.7.1. Overview

The IPPM working group in the IETF defines common criteria and metrics for measuring performance of IP traffic ([IPPM-FW]). Some of the key RFCs published by this working group have defined metrics for measuring connectivity [IPPM-Con], delay ([IPPM-1DM], [IPPM-2DM]), and packet loss [IPPM-1LM]. It should be noted that the work of the IETF in the context of performance metrics is not limited to IP networks; [PM-CONS] presents general guidelines for considering new performance metrics.

The IPPM working group has defined not only metrics for performance measurement, but also protocols that define how the measurement is carried out. The One-way Active Measurement Protocol [OWAMP] and the Two-Way Active Measurement Protocol [TWAMP] define a method and protocol for measuring performance metrics in IP networks.

OWAMP [OWAMP] enables measurement of one-way characteristics of IP networks, such as one-way packet loss and one-way delay. For its proper operation OWAMP requires accurate time of day setting at its end points.

TWAMP [TWAMP] is a similar protocol that enables measurement of both one-way and two-way (round trip) characteristics.

OWAMP and TWAMP are both comprised of two separate protocols:

- o OWAMP-Control/TWAMP-Control: used to initiate, start, and stop test sessions and to fetch their results. Continuity Check and Connectivity Verification are tested and confirmed by establishing the OWAMP/TWAMP Control Protocol TCP connection.
- o OWAMP-Test/TWAMP-Test: used to exchange test packets between two measurement nodes. Enables the loss and delay measurement functions, as well as detection of other anomalies, such as packet duplication and packet reordering.

It should be noted that while [OWAMP] and [TWAMP] define tools for performance measurement, they do not define the accuracy of these tools. The accuracy depends on scale, implementation and network configurations.

Alternative protocols for performance monitoring are defined, for example, in MPLS-TP OAM ([MPLS-LM-DM], [TP-LM-DM]), and in Ethernet OAM [ITU-T-Y1731].

4.7.2. Control and Test Protocols

OWAMP and TWAMP control protocols run over TCP, while the test protocols run over UDP. The purpose of the control protocols is to initiate, start, and stop test sessions, and for OWAMP to fetch results. The test protocols introduce test packets (which contain sequence numbers and timestamps) along the IP path under test according to a schedule, and record statistics of packet arrival. Multiple sessions may be simultaneously defined, each with a session identifier, and defining the number of packets to be sent, the amount of padding to be added (and thus the packet size), the start time, and the send schedule (which can be either a constant time between test packets or exponentially distributed pseudo-random). Statistics recorded conform to the relevant IPPM RFCs.

From a security perspective, OWAMP and TWAMP test packets are hard to detect because they are simply UDP streams between negotiated port numbers, with potentially nothing static in the packets. OWAMP and TWAMP also include optional authentication and encryption for both control and test packets.

4.7.3. OWAMP

OWAMP defines the following logical roles: Session-Sender, Session-Receiver, Server, Control-Client, and Fetch-Client. The Session-Sender originates test traffic that is received by the Session-Receiver. The Server configures and manages the session, as well as returning the results. The Control-Client initiates requests for test sessions, triggers their start, and may trigger their termination. The Fetch-Client requests the results of a completed session. Multiple roles may be combined in a single host - for example, one host may play the roles of Control-Client, Fetch-Client, and Session-Sender, and a second playing the roles of Server and Session-Receiver.

In a typical OWAMP session the Control-Client establishes a TCP connection to port 861 of the Server, which responds with a server greeting message indicating supported security/integrity modes. The Control-Client responds with the chosen communications mode and the Server accepts the mode. The Control-Client then requests and fully describes a test session to which the Server responds with its acceptance and supporting information. More than one test session may be requested with additional messages. The Control-Client then starts a test session and the Server acknowledges, and instructs the Session-Sender to start the test. The Session-Sender then sends test packets with pseudorandom padding to the Session-Receiver until the session is complete or until the Control-client stops the session. Once finished, the Session-Sender reports to the Server which recovers data from the Session-Receiver. The Fetch-Client can then send a fetch request to the Server, which responds with an acknowledgement and immediately thereafter the result data.

4.7.4. TWAMP

TWAMP defines the following logical roles: session-sender, session-reflector, server, and control-client. These are similar to the OWAMP roles, except that the Session-Reflector does not collect any packet information, and there is no need for a Fetch-Client.

In a typical TWAMP session the Control-Client establishes a TCP connection to port 862 of the Server, and mode is negotiated as in OWAMP. The Control-Client then requests sessions and starts them. The Session-Sender sends test packets with pseudorandom padding to the Session-Reflector which returns them with insertion of timestamps.

4.8. TRILL

The requirements of OAM in TRILL are defined in [TRILL-OAM]. The challenge in TRILL OAM, much like in MPLS networks, is that traffic between RBridges RB1 and RB2 may be forwarded through more than one path. Thus, an OAM protocol between RBridges RB1 and RB2 must be able to monitor all the available paths between the two RBridge.

During the writing of this document the detailed definition of the TRILL OAM tools are still work in progress. This subsection presents the main requirements of TRILL OAM.

The main requirements defined in [TRILL-OAM] are:

- o Continuity Checking (CC) - the TRILL OAM protocol must support a function for CC between any two RBridges RB1 and RB2.
- o Connectivity Verification (CV) - connectivity between two RBridges RB1 and RB2 can be verified on a per-flow basis.
- o Path Tracing - allows an RBridge to trace all the available paths to a peer RBridge.
- o Performance monitoring - allows an RBridge to monitor the packet loss and packet delay to a peer RBridge.

5. Summary

This section summarizes the OAM tools and functions presented in this document. This summary is an index to some of the main OAM tools defined in the IETF. This compact index that can be useful to all readers from network operators to standards development organizations. The summary includes a short subsection that presents some guidance to network equipment vendors.

5.1. Summary of OAM Tools

This subsection provides a short summary of each of the OAM toolsets described in this document.

A detailed list of the RFCs related to each toolset is given in Appendix A.1.

+-----+-----+-----+		
Toolset	Description	Transport Technology

IP Ping	Ping ([IntHost], [NetTerms]) is a simple application for testing reachability that uses ICMP Echo messages ([ICMPv4], [ICMPv6]).	IPv4/IPv6
IP Traceroute	Traceroute ([TCPIP-Tools], [NetTools]) is an application that allows users to trace the path between an IP source and an IP destination, i.e., to identify the nodes along the path. If more than one path exists between the source and destination Traceroute traces *a* path. The most common implementation of Traceroute uses UDP probe messages, although there are other implementations that use different probes, such as ICMP or TCP. Paris Traceroute [PARIS] is an extension that attempts to discover all the available paths from A to B by scanning different values of header fields.	IPv4/IPv6
BFD	Bidirectional Forwarding Detection (BFD) is defined in [BFD] as a framework for a lightweight generic OAM tool. The intention is to define a base tool that can be used with various encapsulation types, network environments, and in various medium types.	generic
MPLS OAM	MPLS LSP Ping, as defined in [MPLS-OAM], [MPLS-OAM-FW] and [LSP-Ping], is an OAM tool for point-to-point and point-to-multipoint MPLS LSPs. It includes two main functions: Ping and Traceroute. BFD [BFD-LSP] is an alternative means for detecting MPLS LSP data plane failures.	MPLS
MPLS-TP OAM	MPLS-TP OAM is defined in a set of RFCs.	MPLS-TP

	The OAM requirements for MPLS Transport Profile (MPLS-TP) are defined in [MPLS-TP-OAM]. Each of the tools in the OAM toolset is defined in its own RFC, as specified in Section A.1.	
Pseudowire OAM	The PWE3 OAM architecture defines control channels that support the use of existing IETF OAM tools to be used for a pseudowire (PW). The control channels that are defined in [VCCV] and [PW-G-ACh] may be used in conjunction with ICMP Ping, LSP Ping, and BFD to perform CC and CV functionality. In addition the channels support use of any of the MPLS-TP based OAM tools for completing their respective OAM functionality for a PW.	Pseudowire
OWAMP and TWAMP	The One Way Active Measurement Protocol [OWAMP] and the Two Way Active Measurement Protocols [TWAMP] are two protocols defined in the IP Performance Metrics (IPPM) working group in the IETF. These protocols allow various performance metrics to be measured, such as packet loss, delay and delay variation, duplication and reordering.	IPv4/IPv6
TRILL OAM	The requirements of OAM in TRILL are defined in [TRILL-OAM]. These requirements include continuity checking, connectivity verification, path tracing and performance monitoring. During the writing of this document the detailed definition of the TRILL OAM tools is work in progress.	TRILL

Table 3 Summary of OAM-related IETF Tools

5.2. Summary of OAM Functions

Table 4 summarizes the OAM functions that are supported in each of the toolsets that were analyzed in this section. The columns of this table are the typical OAM functions described in Section 1.3.

Toolset	Continuity Check	Connectivity Verification	Path Discovery	Performance Monitoring	Other Functions
IP Ping	Echo				
IP Traceroute			Traceroute		
BFD	BFD Control / Echo	BFD Control			RDI using BFD Control
MPLS OAM (LSP Ping)		"Ping" mode	"Traceroute" mode		
MPLS-TP OAM	CC	CV/proactive or on-demand	Route Tracing	-LM -DM	-Diagnostic Test -Lock -Alarm Reporting -Client Failure Indication -RDI
Pseudowire OAM	BFD	-BFD -ICMP Ping -LSP-Ping	LSP-Ping		
OWAMP and	- control			-Delay	

	TRAMP	protocol			measur ement -Packet loss measur ement	
+	-----	+	-----	+	-----	+
	TRILL OAM	CC	CV	Path tracing	-Delay measur ement -Packet loss measur ement	
+	-----	+	-----	+	-----	+

Table 4 Summary of the OAM Functionality in IETF OAM Tools

5.3. Guidance to Network Equipment Vendors

As mentioned in Section 1.4. , it is imperative for OAM tools to be capable of testing the actual data plane in as much accuracy as possible. While this guideline may appear obvious, it is worthwhile to emphasize the key importance of enforcing fate-sharing between OAM traffic that monitors the data plane and the data plane traffic it monitors.

6. Security Considerations

OAM is tightly coupled with the stability of the network. A successful attack on an OAM protocol can create a false illusion of non-existent failures, or prevent the detection of actual ones. In both cases the attack may result in denial of service.

Some of the OAM tools presented in this document include security mechanisms that provide integrity protection, thereby preventing attackers from forging or tampering with OAM packets. For example, [BFD] includes an optional authentication mechanism for BFD Control packets, using either SHA1, MD5, or a simple password. [OWAMP] and [TWAMP] have 3 modes of security: unauthenticated, authenticated, and encrypted. The authentication uses SHA1 as the HMAC algorithm, and the encrypted mode uses AES encryption.

Confidentiality is typically not considered a requirement for OAM protocols. However, the use of encryption (e.g., [OWAMP] and

[TWAMP]) can make it difficult for attackers to identify OAM packets, thus making it more difficult to attack the OAM protocol.

OAM can also be used as a means for network reconnaissance; information about addresses, port numbers and about the network topology and performance can be gathered either by passively eavesdropping to OAM packets, or by actively sending OAM packets and gathering information from the respective responses. This information can then be used maliciously to attack the network. Note that some of this information, e.g., addresses and port numbers, can be gathered even when encryption is used ([OWAMP], [TWAMP]).

For further details about the security considerations of each OAM protocol, the reader is encouraged to review the Security Considerations section of each document referenced by this memo.

7. IANA Considerations

There are no new IANA considerations implied by this document.

8. Acknowledgments

The authors gratefully acknowledge Sasha Vainshtein, Carlos Pignataro, David Harrington, Dan Romascanu, Ron Bonica, Benoit Claise, Stewart Bryant, Tom Nadeau, Elwyn Davies, Al Morton, Sam Aldrin, Thomas Narten, and other members of the OPSA WG for their helpful comments on the mailing list.

This document was prepared using 2-Word-v2.0.template.dot.

9. References

9.1. Normative References

[OAM-Def] Andersson, L., Van Helvoort, H., Bonica, R., Romascanu, D., Mansfield, S., "Guidelines for the use of the OAM acronym in the IETF ", RFC 6291, June 2011.

9.2. Informative References

[ATM-L2] Singh, S., Townsley, M., and C. Pignataro, "Asynchronous Transfer Mode (ATM) over Layer 2 Tunneling Protocol Version 3 (L2TPv3)", RFC 4454, May 2006.

[BFD] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

- [BFD-Gen] Katz, D., Ward, D., "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, June 2010.
- [BFD-IP] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.
- [BFD-LSP] Aggarwal, R., Kompella, K., Nadeau, T., and Swallow, G., "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.
- [BFD-Multi] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, June 2010.
- [BFD-VCCV] Nadeau, T., Pignataro, C., "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010.
- [Comp] Bonaventure, O., "Computer Networking: Principles, Protocols and Practice", 2008.
- [Dup] Uijterwaal, H., "A One-Way Packet Duplication Metric", RFC 5560, May 2009.
- [Eth-Int] Mohan, D., Bitar, N., Sajassi, A., Delord, S., Niger, P., Qiu, R., "MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking", RFC 7023, October 2013.
- [G-ACh] Bocci, M., Vigoureux, M., Bryant, S., "MPLS Generic Associated Channel", RFC 5586, June 2009.
- [ICMP-Ext] Bonica, R., Gan, D., Tappan, D., Pignataro, C., "ICMP Extensions for Multiprotocol Label Switching", RFC 4950, August 2007.
- [ICMP-Int] Atlas, A., Bonica, R., Pignataro, C., Shen, N., Rivers, JR., "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010.
- [ICMP-MP] Bonica, R., Gan, D., Tappan, D., Pignataro, C., "Extended ICMP to Support Multi-Part Messages", RFC 4884, April 2007.

- [ICMPv4] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.
- [ICMPv6] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [IEEE802.1Q] IEEE 802.1Q, "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", October 2012.
- [IEEE802.3ah] IEEE 802.3, "IEEE Standard for Information technology - Local and metropolitan area networks - Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications", clause 57, December 2008.
- [IntHost] Braden, R., "Requirements for Internet Hosts -- Communication Layers", RFC 1122, October 1989.
- [IPPM-1DM] Almes, G., Kalidindi, S., Zekauskas, M., "A One-way Delay Metric for IPPM", RFC 2679, September 1999.
- [IPPM-1LM] Almes, G., Kalidindi, S., Zekauskas, M., "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.
- [IPPM-2DM] Almes, G., Kalidindi, S., Zekauskas, M., "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.
- [IPPM-Con] Mahdavi, J., Paxson, V., "IPPM Metrics for Measuring Connectivity", RFC 2678, September 1999.
- [IPPM-FW] Paxson, V., Almes, G., Mahdavi, J., and Mathis, M., "Framework for IP Performance Metrics", RFC 2330, May 1998.
- [ITU-G8113.1] ITU-T Recommendation G.8113.1/Y.1372.1, "Operations, Administration and Maintenance mechanism for MPLS-TP in Packet Transport Network (PTN)", November 2012.
- [ITU-G8113.2] ITU-T Recommendation G.8113.2/Y.1372.2, "Operations, administration and maintenance mechanisms for MPLS-TP networks using the tools defined for MPLS", November 2012.

- [ITU-T-CT] Betts, M., "Allocation of a Generic Associated Channel Type for ITU-T MPLS Transport Profile Operation, Maintenance, and Administration (MPLS-TP OAM)", RFC 6671, November 2012.
- [ITU-T-G.806] ITU-T Recommendation G.806, "Characteristics of transport equipment - Description methodology and generic functionality", January 2009.
- [ITU-T-Y1711] ITU-T Recommendation Y.1711, "Operation & Maintenance mechanism for MPLS networks", February 2004.
- [ITU-T-Y1731] ITU-T Recommendation G.8013/Y.1731, "OAM Functions and Mechanisms for Ethernet-based Networks", July 2011.
- [ITU-Terms] ITU-R/ITU-T Terms and Definitions, online, 2013.
- [L2TP-EC] McGill, N. and C. Pignataro, "Layer 2 Tunneling Protocol Version 3 (L2TPv3) Extended Circuit Status Values", RFC 5641, August 2009.
- [L2VPN-OAM] Sajassi, A., Mohan, D., "Layer 2 Virtual Private Network (L2VPN) Operations, Administration, and Maintenance (OAM) Requirements and Framework", RFC 6136, March 2011.
- [L3VPN-OAM] El Mghazli, Y., Nadeau, T., Boucadair, M., Chan, K., Gonguet, A., "Framework for Layer 3 Virtual Private Networks (L3VPN) Operations and Management", RFC 4176, October 2005.
- [Lock-Loop] Boutros, S., Sivabalan, S., Aggarwal, R., Vigoureux, M., Dai, X., "MPLS Transport Profile Lock Instruct and Loopback Functions", RFC 6435, November 2011.
- [LSP-Ping] Kompella, K., Swallow, G., "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [Mng] Farrel, A., "Inclusion of Manageability Sections in Path Computation Element (PCE) Working Group Drafts", RFC 6123, February 2011.
- [MPLS-ENCAPS] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T. and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.

- [MPLS-LM-DM] Frost, D., Bryant, S., "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.
- [MPLS-OAM] Nadeau, T., Morrow, M., Swallow, G., Allan, D., Matsushima, S., "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.
- [MPLS-OAM-FW] Allan, D., Nadeau, T., "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", RFC 4378, February 2006.
- [MPLS-P2MP] Yasukawa, S., Farrel, A., King, D., Nadeau, T., "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks", RFC 4687, September 2006.
- [MPLS-TP-OAM] Vigoureux, M., Ward, D., Betts, M., "Requirements for OAM in MPLS Transport Networks", RFC 5860, May 2010.
- [mtrace] Fenner, W., Casner, S., "A "traceroute" facility for IP Multicast", draft-ietf-idmr-traceroute-ipm-07 (expired), July 2000.
- [NetTerms] Jacobsen, O., Lynch, D., "A Glossary of Networking Terms", RFC 1208, March 1991.
- [NetTools] Enger, R., Reynolds, J., "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices", RFC 1470, June 1993.
- [OAM-Analys] Sprecher, N., Fang, L., "An Overview of the OAM Tool Set for MPLS based Transport Networks", RFC 6669, July 2012.
- [OAM-Label] Ohta, H., "Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions", RFC 3429, November 2002.
- [OAM-Mng] Ersue, M., Claise, B., "An Overview of the IETF Network Management Standards", RFC 6632, June 2012.

- [OnDemand-CV] Gray, E., Bahadur, N., Boutros, S., Aggarwal, R. "MPLS On-Demand Connectivity Verification and Route Tracing", RFC 6426, November 2011.
- [OWAMP] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and Zekauskas, M., "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.
- [PARIS] Brice Augustin, Timur Friedman and Renata Teixeira, "Measuring Load-balanced Paths in the Internet", IMC, 2007.
- [PM-CONS] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, October 2011.
- [PW-ACH] Bryant, S., Swallow, G., Martini, L., McPherson, D., "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [PW-G-ACh] Li, H., Martini, L., He, J., Huang, F., "Using the Generic Associated Channel Label for Pseudowire in the MPLS Transport Profile (MPLS-TP)", RFC 6423, November 2011.
- [PW-MAP] Aissaoui, M., Busschbach, P., Martini, L., Morrow, M., Nadeau, T., and Y(J). Stein, "Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping", RFC 6310, July 2011.
- [Reorder] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, November 2006.
- [Signal] Yasukawa, S., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.
- [TCPIP-Tools] Kessler, G., Shepard, S., "A Primer On Internet and TCP/IP Tools and Utilities", RFC 2151, June 1997.
- [TP-CC-CV] Allan, D., Swallow, G., Drake, J., "Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile", RFC 6428, November 2011.

- [TP-Fault] Swallow, G., Fulignoli, A., Vigoureux, M., Boutros, S., "MPLS Fault Management Operations, Administration, and Maintenance (OAM)", RFC 6427, November 2011.
- [TP-LM-DM] Frost, D., Bryant, S., "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, September 2011.
- [TP-OAM-FW] Busi, I., Allan, D., "Operations, Administration and Maintenance Framework for MPLS-based Transport Networks ", RFC 6371, September 2011.
- [TP-Term] Van Helvoort, H., Andersson, L., Sprecher, N., "A Thesaurus for the Terminology used in MPLS Transport Profile (MPLS-TP) Internet-Drafts and RFCs in the Context of the ITU-T's Transport Network Recommendations", RFC 7087, December 2013.
- [TRILL-OAM] Senevirathne, T., Bond, D., Aldrin, S., Li, Y., Watve, R., "Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)", RFC 6905, March 2013.
- [TWAMP] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and Babiarz, J., "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.
- [VCCV] Nadeau, T., Pignataro, C., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.
- [VCCV-SURVEY] Del Regno, N., Malis, A., "The Pseudowire (PW) and Virtual Circuit Connectivity Verification (VCCV) Implementation Survey Results", RFC 7079, November 2013.

Appendix A.

List of OAM Documents

A.1. List of IETF OAM Documents

Table 5 summarizes the OAM related RFCs published by the IETF.

It is important to note that the table lists various RFCs that are different by nature. For example, some of these documents define OAM tools or OAM protocols (or both), while others define protocols that

are not strictly OAM-related, but are used by OAM tools. The table also includes RFCs that define the requirements or the framework of OAM in a specific context (e.g., MPLS-TP).

The RFCs in the table are categorized in a few sets as defined in Section 1.3.

Toolset	Title	RFC
IP Ping	Requirements for Internet Hosts -- Communication Layers [IntHost]	RFC 1122
	A Glossary of Networking Terms [NetTerms]	RFC 1208
	Internet Control Message Protocol [ICMPv4]	RFC 792
	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification [ICMPv6]	RFC 4443
IP Traceroute	A Primer On Internet and TCP/IP Tools and Utilities [TCPIP-Tools]	RFC 2151
	FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices [NetTools]	RFC 1470
	Internet Control Message Protocol [ICMPv4]	RFC 792
	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification [ICMPv6]	RFC 4443
	Extended ICMP to Support Multi-Part Messages [ICMP-MP]	RFC 4884

	Extending ICMP for Interface and Next-Hop Identification [ICMP-Int]	RFC 5837
BFD	Bidirectional Forwarding Detection [BFD]	RFC 5880
	Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop) [BFD-IP]	RFC 5881
	Generic Application of Bidirectional Forwarding Detection [BFD-Gen]	RFC 5882
	Bidirectional Forwarding Detection (BFD) for Multihop Paths [BFD-Multi]	RFC 5883
	Bidirectional Forwarding Detection for MPLS Label Switched Paths (LSPs) [BFD-LSP]	RFC 5884
	Bidirectional Forwarding Detection for the Pseudowire Virtual Circuit Connectivity Verification (VCCV) [BFD-VCCV]	RFC 5885
MPLS OAM	Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks [MPLS-OAM]	RFC 4377
	A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM) [MPLS-OAM-FW]	RFC 4378
	Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures [LSP-Ping]	RFC 4379
	Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks [MPLS-P2MP]	RFC 4687

	ICMP Extensions for Multiprotocol Label Switching [ICMP-Ext]	RFC 4950
	Bidirectional Forwarding Detection for MPLS Label Switched Paths (LSPs) [BFD-LSP]	RFC 5884
MPLS-TP OAM	Requirements for OAM in MPLS-TP [MPLS-TP-OAM]	RFC 5860
	MPLS Generic Associated Channel [G-ACh]	RFC 5586
	MPLS-TP OAM Framework [TP-OAM-FW]	RFC 6371
	Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile [TP-CC-CV]	RFC 6428
	MPLS On-Demand Connectivity Verification and Route Tracing [OnDemand-CV]	RFC 6426
	MPLS Fault Management Operations, Administration, and Maintenance (OAM) [TP-Fault]	RFC 6427
	MPLS Transport Profile Lock Instruct and Loopback Functions [Lock-Loop]	RFC 6435
	Packet Loss and Delay Measurement for MPLS Networks [MPLS-LM-DM]	RFC 6374
	A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks [TP-LM-DM]	RFC 6375
Pseudowire	Pseudowire Virtual Circuit	RFC 5085

OAM	Connectivity Verification (VCCV): A Control Channel for Pseudowires [VCCV]	
	Bidirectional Forwarding Detection for the Pseudowire Virtual Circuit Connectivity Verification (VCCV) [BFD-VCCV]	RFC 5885
	Using the Generic Associated Channel Label for Pseudowire in the MPLS Transport Profile (MPLS-TP) [PW-G-ACh]	RFC 6423
	Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping [PW-MAP]	RFC 6310
	MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking [Eth-Int]	RFC 7023
OWAMP and TWAMP	A One-way Active Measurement Protocol [OWAMP]	RFC 4656
	A Two-Way Active Measurement Protocol [TWAMP]	RFC 5357
	Framework for IP Performance Metrics [IPPM-FW]	RFC 2330
	IPPM Metrics for Measuring Connectivity [IPPM-Con]	RFC 2678
	A One-way Delay Metric for IPPM [IPPM-1DM]	RFC 2679
	A One-way Packet Loss Metric for IPPM [IPPM-1LM]	RFC 2680
	A Round-trip Delay Metric for IPPM	RFC 2681

	[IPPM-2DM]	
	Packet Reordering Metrics [Reorder]	RFC 4737
	A One-Way Packet Duplication Metric [Dup]	RFC 5560
TRILL OAM	Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)	RFC 6905

Table 5 Summary of IETF OAM Related RFCs

A.2. List of Selected Non-IETF OAM Documents

In addition to the OAM tools defined by the IETF, the IEEE and ITU-T have also defined various OAM tools that focus on Ethernet, and various other transport network environments. These various tools, defined by the three standard organizations, are often tightly coupled, and have had a mutual effect on each other. The ITU-T and IETF have both defined OAM tools for MPLS LSPs, [ITU-T-Y1711] and [LSP-Ping]. The following OAM standards by the IEEE and ITU-T are to some extent linked to IETF OAM tools listed above and are mentioned here only as reference material:

- o OAM tools for Layer 2 have been defined by the ITU-T in [ITU-T-Y1731], and by the IEEE in 802.1ag [IEEE802.1Q] . The IEEE 802.3 standard defines OAM for one-hop Ethernet links [IEEE802.3ah].
- o The ITU-T has defined OAM for MPLS LSPs in [ITU-T-Y1711], and MPLS-TP OAM in [ITU-G8113.1] and [ITU-G8113.2].

It should be noted that these non-IETF documents deal in many cases with OAM functions below the IP layer (Layer 2, Layer 2.5) and in some cases operators use a multi-layered OAM approach, which is a function of the way their networks are designed.

Table 6 summarizes some of the main OAM standards published by non-IETF standard organizations. This document focuses on IETF OAM standards, but these non-IETF standards are referenced in this document where relevant.

	Title	Standard/Draft
ITU-T MPLS OAM	Operation & Maintenance mechanism for MPLS networks [ITU-T-Y1711]	ITU-T Y.1711
	<p>Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions [OAM-Label]</p> <p>Note: although this is an IETF document, it is listed as one of the non-IETF OAM standards, since it was defined as a complementary part of ITU-T Y.1711.</p>	RFC 3429
ITU-T MPLS-TP OAM	<p>Operations, administration and Maintenance mechanisms for MPLS-TP networks using the tools defined for MPLS [ITU-G8113.2]</p> <p>Note: this document describes the OAM toolset defined by the IETF for MPLS-TP, whereas ITU-T G.8113.1 describes the OAM toolset defined by the ITU-T.</p>	ITU-T G.8113.2
	Operations, Administration and Maintenance mechanism for MPLS-TP in Packet Transport Network (PTN)	ITU-T G.8113.1
	<p>Allocation of a Generic Associated Channel Type for ITU-T MPLS Transport Profile Operation, Maintenance, and Administration (MPLS-TP OAM) [ITU-T-CT]</p> <p>Note: although this is an IETF document, it is listed as one of the</p>	RFC 6671

	non-IETF OAM standards, since it was defined as a complementary part of ITU-T G.8113.1.	
ITU-T Ethernet OAM	OAM Functions and Mechanisms for Ethernet-based Networks [ITU-T-Y1731]	ITU-T Y.1731
IEEE CFM	Connectivity Fault Management [IEEE802.1Q] Note: CFM was originally published as IEEE 802.1ag, but is now incorporated in the 802.1Q standard.	IEEE 802.1ag
IEEE DDCFM	Management of Data Driven and Data Dependent Connectivity Faults [IEEE802.1Q] Note: DDCFM was originally published as IEEE 802.1Qaw, but is now incorporated in the 802.1Q standard.	IEEE 802.1ag
IEEE 802.3 link level OAM	Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks [IEEE802.3ah] Note: link level OAM was originally defined in IEEE 802.3ah, and is now incorporated in the 802.3 standard.	IEEE 802.3ah

Table 6 Non-IETF OAM Standards Mentioned in this Document

Authors' Addresses

Tal Mizrahi
Marvell
6 Hamada St.
Yokneam, 20692
Israel

Email: talmi@marvell.com

Nurit Sprecher
Nokia Solutions and Networks
3 Hanagar St. Neve Ne'eman B
Hod Hasharon, 45241
Israel

Email: nurit.sprecher@nsn.com

Elisa Bellagamba
Ericsson
6 Farogatan St.
Stockholm, 164 40
Sweden

Phone: +46 761440785
Email: elisa.bellagamba@ericsson.com

Yaacov Weingarten
34 Hagefen St.
Karnei Shomron, 4485500
Israel

Email: wyaacov@gmail.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 16, 2011

B. Nordman
Lawrence Berkeley National
Laboratory
R. Winter
NEC Labs Europe
March 15, 2011

Considerations for Power and Energy Management
draft-norwin-energy-consider-02

Abstract

With rising cost and an increasing awareness of the environmental impact of energy consumption, a desirable feature of networked devices is to be able to assess their power state and energy consumption at will. With this data available, one can build sophisticated applications such as monitoring applications or even active energy management systems. These systems themselves are out of scope of this memo, as it discusses only considerations for the monitored devices. Implementation specifics such as the definition of a Management Information Base are also outside the scope of this document.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 16, 2011.

Copyright Notice

Copyright (c) 2011 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.

Table of Contents

1. Requirements notation	3
2. Overview/Goals	4
3. Settled topics	5
3.1. Scope of Devices	5
3.2. Identity	5
3.3. Power Levels	5
3.4. Devices	5
3.5. Intervals	6
3.6. Presentation to non-IETF audiences	6
3.7. Functions vs. Entities	6
3.8. Simple and Complex Devices	6
4. Topics under discussion	7
4.1. Power States	7
5. Energy Manangement	8
5.1. Control	8
5.2. Identity	8
5.3. NMS Considerations	9
5.4. MIB Considerations	9
5.5. Power Considerations	10
5.6. Incomplete data	10
5.7. Time reporting	10
5.8. Portable devices	10
5.9. Beyond energy	11
5.10. Power State Monitoring	11
5.11. Power Distribution	11
6. Use Context and Use Cases	12
7. Future Directions	13
8. Security Considerations	14
9. Normative References	15
Authors' Addresses	16

1. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Overview/Goals

This document aims at framing discussions on power and energy management within the IETF and recording their results. It clarifies terminology that is routinely used to have multiple contrary meanings, which results in unnecessary confusion. The document further describes how energy and power reporting differs from other reporting tasks that have been defined by the IETF and the resulting implications for mechanisms the IETF will define. This document is intended to be a living document that also captures why certain decisions were made in the process of defining power and energy management mechanisms.

3. Settled topics

The following are topics that seem settled in eman discussions, recognizing that this draft has no authority on that point.

3.1. Scope of Devices

All energy-using devices that have a network connection are in scope. The eman mechanisms also provide for non-IP devices that are supplied with power or that have power metered by an IP device, or are brought into the eman context by a gateway/proxy.

While first adopters will surely be devices such as switches, routers, and servers (some of which already report power levels and power state through proprietary means), in the future networked electronic devices, appliances, and even lights will also need such capability. These devices may have different ways of accomplishing discovery and management for functional purposes, but will share the common energy and power reporting capability. While some devices will directly measure power, other devices will not be able to measure their power, but may be able to reliably estimate it. These devices are still in scope.

3.2. Identity

Some universal mechanisms for identity are needed so that the NMS knows what the devices are that are using energy. The nature of these mechanisms, whether they are existing ones to be referenced or new ones to be created (almost certainly some of both) has not yet been determined.

3.3. Power Levels

The power level of a device is its current electricity demand. It is an important complement to power mode, providing articulation of power level within the basic mode. It also avoids the need for a large number of named modes. Basic modes are distinguished by important functional differences or power levels. Core power modes are an abstraction from individual implementations.

3.4. Devices

The organizing unit for power is a single device with one or more power sources. The term "product" is sometimes used as a synonym, and also covers the case in which a device proxies network presence including power reporting for a second device.

3.5. Intervals

A common feature of energy monitoring is to track energy use over time. Recording of energy use for intervals of time is the responsibility of a network management system (or whatever entity requests data via the eman protocol), not the monitored device itself. The monitored device always reports accumulated energy use with an associated timestamp.

3.6. Presentation to non-IETF audiences

Many people and organizations who have not in the past understood or interacted with the IETF will be interested in eman results. They need to be provided with easily understandable explanations of what eman does and why. How this presentation will be accomplished is still to be determined.

3.7. Functions vs. Entities

Eman is concerned with exposing information to Network Management Systems (NMSs). Providing information is a function. The various functions may be implemented by a single device, or distributed among several devices.

3.8. Simple and Complex Devices

We will support both. Simple devices want to avoid complexity that burdens both implementation on the monitored device, and the monitoring system. Complex devices need to have access to additional data fields and capabilities.

4. Topics under discussion

4.1. Power States

We synonymously use the terms Power Mode and Power State; named modes are general categories only ("buckets"), not individual states with highly-specific meaning.

Discussions about energy consumptions and device power states are often confusing as different products define states such as "standby" quite differently. Even the same class of devices often implement named states differently. Named power states are intrinsically difficult to define consistently as they imply not only something about a device's energy consumption but also something about the device's capabilities in that state, and are implementation-dependent. All of this makes highly-specific named modes unsuitable for use in a general context. The term with by far the most different definitions is "standby" and so we therefore do not refer to standby in this document and believe it unsuitable for use in eman.

We believe that the three named power state categories, on, off and sleep, are broadly understood. These mode categories may each contain a large set of power sub-states. A fourth basic power state of 'ready' may be more appropriate for some devices, particularly appliances.

In general, devices that are asleep will be able to wake quickly and will retain network connectivity. Devices that are off usually take much more time to turn on than the wake time and usually lack network connectivity. Devices that are on are fully functional but potentially with reduced performance.

A critical feature of the set of basic power states is that they should be universally applicable to any device eman is applied to. This does not mean that each device has every state, but that the model is sufficiently general that it can be applied to all. When the level of detail rises, the set of states usually is then applicable to only certain types of products, and/or to specific implementations. In addition, these detailed states generally embody specific functional characteristics of the state, and so are better embodied in other variables (that may be delivered by an energy management protocol).

5. Energy Manangement

First and foremost, the task of power and energy management is reporting. While a more active role in energy management is conceivable by e.g. putting devices into power states based on policies or other predefined schemes at a network management system (NMS).

5.1. Control

There should not be an assumption that power state management of devices is done externally/centrally. Ideally most devices will manage their own power state, implementing distributed intelligence. The control function is accomplished separately from power reporting. A core mechanism many devices will use to manage power consumption is a price (and price forecast) for electricity.

5.2. Identity

All devices on a network need to expose identity to others. While some protocols accomplish this for particular applications or contexts, it is desirable to have a simple universal mechanism. This is particularly true for devices that may have a fairly limited degree of participation in the network, such as appliances.

For energy management purposes, the it is important to know "what" a device is, and "who" it is. Each of these has two parts as follows:

- o "Species". This is the fundamental classification that a device is a member of due to its design and capabilities. This property is determined by the manufacturer before it is sold. Examples are server, router, notebook PC, display, TV, refrigerator, light, etc.
- o "Origin". The brand and model of the device. Primarily a method to find out more information about a device, such as its specifications for requirements and capabilities. It would be advantageous to include a URL for detailed information from the manufacturer. An example of this is the "Universal Product Code" on many products.
- o Name: A human-readable name, locally specified when the device is configured or installed.
- o Network ID: A globally unique identifier for the NMS to use to recognize a device. This should be based on one or more existing IETF mechanisms.

An energy management application could then obtain current energy use for a device like a refrigerator, and compare it to what it is expected to use under normal operation, and alert the building manager if it is significantly out of range. This also can be used to quickly inventory energy-using products in a building, and to summarize by product type where energy is being used.

5.3. NMS Considerations

A Network Management System is an entity which collects energy and power reporting data and uses it for advanced applications. One such application correlates energy consumption with other metrics to display efficiency metrics (like watthours/bit). An NMS can also set device policies to control larger networked systems such as a data center.

An NMS will query energy MIB data on a periodic basis, with that period dictated by its needs, possibly being dynamic. MIBs should provide an energy "meter reading" to allow computing of energy use for any period. Thus, the NMS does most of the work to generate time series energy data, and this minimizes burden on the host and the complexity of the Power MIB.

The core function of power monitoring is to maintain meters of energy use and of time in different power states (and through summing, total energy and time). The second is to be able to report current power consumption and power state.

5.4. MIB Considerations

The MIB should be generic as there are a large number of devices yet to come and power states are and will become more diverse.

The MIB should be structured so that the smallest possible set of values/information is applicable to a large range of devices, can be implemented efficiently and is extensible to accommodate additional information objects. As an example, many devices will not be battery powered but it should be easy to add battery monitoring to the basic set of energy-related information.

The proposed MIB structures enable reporting on components of products (e.g. linecards in a chassis) in addition to entire products. Doing this is not part of the eman charter, so while there is no reason to preclude the capability, it should not be a distraction to completing the chartered eman scope.

5.5. Power Considerations

Reporting should cover both AC and DC power sources. However, other types should be provided for, and the type of energy is one of the reported values. Standard low-voltage DC (e.g. USB, Power over Ethernet, eMerge) is immediately useful. A core set of values should be available from any device that implements the Power MIB at all so that an NMS can quickly obtain and aggregate uniform data for all devices.

There is a fundamental distinction between supplied power from a device And input power to a device, notably losses that occur in transmission, as well as other (possibly unknown) devices that are also using the power. The effect of internal batteries is not revealed by the MIB, as it only reports on net power into or out of a device.

5.6. Incomplete data

Energy reporting will cover a wide variety of information about a device, its status, and energy usage. Sometimes, particularly for legacy or non-IP products, this will be incomplete. It is critical that the fact that some data are missing does not undermine the ability to report the data that are present.

5.7. Time reporting

At the core of energy reporting is data from energy meters that are meter readings associated with timestamps. A variety of issues arise on the meaning of that time.

Without strong synchronization, the NMS and the devices it queries will have different absolute times. However, the NMS knows when it asked for each meter reading so can account for this difference.

For some devices, when they are off they will be unable to accumulate their energy consumption. The fact that some consumption may be missing needs to be communicated to the NMS. One possibility is to record the last time that a period of missing energy occurred, and report that to the NMS.

5.8. Portable devices

Devices that are routinely moved from one building to another (or even within a building) pose special challenges for energy reporting. The question arises whether it is the energy into the device, or from the building, which is dominant. It may be important to record the time a device most recently changed power domain to ensure that a NMS

can correctly account only for energy consumed on its premises.

5.9. Beyond energy

The charter references "energy" but virtually all discussion has been limited to electricity. Other forms of energy should be included at some point; we should discuss whether this is readily feasible now, or needs to be postponed to future work.

5.10. Power State Monitoring

For the device power state, the following information is considered to be relevant:

- o the current state
- o the time of (or time since) the last change
- o the current real power (energy consumption rate)
- o accumulated energy consumption

5.11. Power Distribution

Wired networks enable power distribution that is co-incident with network Communication. However, many devices will not communicate on the same Medium that they are powered on, or may lack connectivity entirely (though with the power provider knowing of their identity). Devices can report power for another device only if they are the entity providing the power.

6. Use Context and Use Cases

The following are some use contexts that this facility is intended for. These are not necessarily mutually exclusive, and a device can report the same data regardless of the context.

- o A data center, with a NMS which is integrated with application functionality, and also manages energy use.
- o A commercial building, in which the energy reporting is separate from any management of devices, and more as background to help understand building operation (including occupancy) and identify inefficiencies or equipment failures.
- o A house, which shares some of the commercial building characteristics, but with different management approach and security concerns.
- o A vehicle, which uses the reporting only for automatic management, not for reporting to the user.

Use cases include a facility manager or an NMS in an automated fashion:

- o Understand costs for billing purposes.
- o Assess savings potentials.
- o Identify possible device malfunctions.
- o Reveal unexpected usage patterns.
- o Plan for future capacity needs.
- o Understand heat production in a building or space.
- o A NMS which deals with draws on current power use to deal with an actual or potential shortfall in power supply.

7. Future Directions

The current effort to create a protocol for energy management is unlikely to be the last word on the topic. In fact, there are many directions that need to be explored for potential addition to the features enabled by this mechanism or others. These include:

- o other energy media such as wireless power, non-electric energy (e.g. natural gas, steam, hot/cold water).
- o more features for control.
- o other energy-relevant quantities (e.g. temperatures, flow rates).
- o other resources (e.g. water).

8. Security Considerations

None.

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Authors' Addresses

Bruce Nordman
Lawrence Berkeley National Laboratory

Email: bnordman@lbl.gov

Rolf Winter
NEC Labs Europe

Email: rolf.winter@neclab.eu

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 28, 2011

J. Quittek, Ed.
R. Winter
T. Dietz
D. Dudkowski
NEC Europe Ltd.
October 25, 2010

Definition of Managed Objects for Energy Management
draft-quittek-power-mib-02.txt

Abstract

This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes managed objects providing information about the energy consumption, the power states, and the battery status of managed devices and their components.

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/lid-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 April 28, 2011.

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 BSD License.

Table of Contents

1. The Internet-Standard Management Framework	3
2. Introduction	3
3. Identifying Monitored Devices and Components	4
4. Power State MIB	5
4.1. Current Power State Table	5
4.2. Power State Table	6
5. Energy MIB	7
5.1. Energy Consumption Table	7
5.2. Energy Consumption Per Power State Table	8
5.3. Power History Table	9
6. Battery MIB	9
7. Relationship to Other MIB Modules	10
8. Definitions	11
8.1. Power State MIB	11
8.2. Energy MIB	20
8.3. Battery MIB	39
9. Security Considerations	51
10. IANA Considerations	52
11. References	52
11.1. Normative References	52
11.2. Informative References	53
Authors' Addresses	54

1. The Internet-Standard Management Framework

For a detailed overview of the documents that describe the current Internet-Standard Management Framework, please refer to section 7 of RFC 3410 [RFC3410].

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB are defined using the mechanisms defined in the Structure of Management Information (SMI). This memo specifies MIB modules that is compliant to the SMIV2, which is described in STD 58, RFC 2578 [RFC2578], STD 58, RFC 2579 [RFC2579] and STD 58, RFC 2580 [RFC2580].

2. Introduction

Energy management in communication networks is a topic that has been neglected for many years when energy was cheap and global warming not recognized. This has changed recently. Energy management is becoming a significant component of network planning, operations and management and new energy management strategies are currently being explored.

An essential requirement for energy management is collecting information on energy consumption and energy storage at managed devices.

An elementary step into this direction is monitoring power states. A power state defines a limitations of services provided by a device and implicitly limits energy consumption. Examples for commonly implemented power states include 'on', 'full power', 'low power', 'sleep', 'stand-by', and 'off'. There is no commonly agreed convention for power states naming and semantics. Therefore power states with the same names may have different semantics and different names may be in use for the same power state.

But the actual energy consumption of a device depends on more than just its power state. Also the current load, the kind of load, and many other factors influence energy consumption. If instrumentation is available, it is very helpful to receive information on the actual energy consumption of a device and its component. Providing this information requires much more effort than reporting power states, because a probe that measures (electrical) power is required. Typically this means not just adding several lines of software to a device, but also adding costly sensor hardware to it.

A third aspect to be considered for energy management is energy storage in batteries. It is helpful, for example, to monitor which device is running on batteries and which is charging its battery. Fortunately, the problem of instrumentation is often an easy one for devices with rechargeable batteries. Controlling the charging cycles needs instrumentation anyway and this instrumentation can also be used for providing battery status information.

This document defines a portion of the Management Information Base (MIB) that serves the three purposes sketched above:

- o monitoring power states of managed entities,
- o monitoring energy consumption of managed entities,
- o monitoring the status of batteries contained in or controlled by managed devices.

Supporting all three monitoring task will not make sense for every device. Many networked devices do not have batteries to be monitored and thus it would not make sense for them to implement managed objects for this purpose.

As mentioned above, instrumentation for measuring actual energy consumption is relatively expensive and it will not make sense for every managed device to provide sufficient instrumentation. In such a case it would not be appropriate to still implement managed objects for energy consumption monitoring.

This leads to the conclusion that the portions of the MIB for the three monitoring tasks listed above should be rather independent of each other and not combined in a single one. This document contains three MIB modules called Power State MIB, Energy MIB, and Battery MIB. The Energy MIB module uses an object defined in the Power State MIB module, but beyond that there is no dependency between the three modules. Obviously, any combination of the three modules is possible.

The definitions in this document are based on the requirements outlined in [I-D.quittek-power-monitoring-requirements].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

3. Identifying Monitored Devices and Components

As argued in [I-D.quittek-power-monitoring-requirements] it is often required or at least desirable to not just monitor energy consumption and power state of an entire devices, but also of its contained

individual components. Furthermore it is argued in [I-D.quittek-power-monitoring-requirements] that there are cases where it is required that a managed device reports about energy consumption of one or more other, potentially remote devices. An example is a power strip reporting actual power and accumulated energy consumption of devices plugged into it.

It is not the purpose of MIB modules in this document to solve the problem of identifying components of the managed device that implements these modules or of components remote to this managed device. The task of identifying the entity that is subject of monitoring is left to other MIB modules, such as the ENERGY AWARE MIB module [I-D.pareello-eman-energy-aware-mib], and the Entity MIB module [RFC4133].

As an open and flexible way of identifying the monitored entity, the MIB modules in this document use an OID as index that points into a MIB module used for identifying the monitored entity. For simplifying the trivial case that the monitored entity is identical with the device that implements the MIB module, an empty OID may be used.

4. Power State MIB

A number of devices today can operate in a number of different power states by reducing performance or going into standby mode or sleep mode. The Power State MIB module can be used for monitoring these states. Typically, not much instrumentation is needed for supporting the power state MIB module, because most devices with different power states are already equipped with means for controlling their these.

The Power State MIB module is structured into two tables, the powerCurrentStateTable reporting the current power state per entity and the powerStateTable providing statistics per power state. In addition, the Power State MIB module defines a notification that can be sent for informing the receiver about a change of an entity's current power state. For identifying the entity for which power state information is provided, OIDs are used, as explained in the previous section. Both tables use such an OID as their first index.

4.1. Current Power State Table

For basic monitoring of the actual power state of an entity, there is already a MIB module available: the Entity State MIB [RFC4268]. It reports the power state of an entity in object entStateStandby. It can have four different values: unknown(0), off(1), nonOperational(2), operational(3), see ENTITY-STATE-TC-MIB in

[RFC4268].

If this was considered to be sufficient, there would be no need for replicating this object in the power state MIB module. However, there is a concern that the three "known" states are too few for reflecting the variety of power saving states available today. For PCs, for example, there are several more states defined for the Advanced Configuration & Power Interface (ACPI). It might be useful to support several or all of these power states as suggested by [I-D.claise-energy-monitoring-mib].

The powerCurrentStateTable contains just a two objects per row:

```
powerStateTable(1)
+--powerStateEntry(1) [powerStateEnergyConsumerId]
  +-- --- Integer32      powerStateEnergyConsumerId(1)
  +-- --- ObjectIdentifier powerStateEnergyConsumerOid(2)
  +-- r-n SnmpAdminString powerStateOperationalState(3)
  +-- rwn SnmpAdminString powerStateAdminState(4)
```

Object powerStateOperationalState reports the actual power state of an entity at the time the object's value is retrieved. Object powerStateAdminState indicates a desired power state that the entity has been requested to enter, for example, by a network management system.

4.2. Power State Table

The second table called powerStateTable provides more detailed statistics for each power state. For this purpose it uses the power state name as another index object next to the entity index. This way, statistics can be reported per entity and per power state. The second index has the syntax of a SnmpAdminString and can be defined by the manufacturer of the device or MIB. In this way the index can fit many devices because the characteristics of the power state can be defined per device. The characteristics of the power state SHOULD be described as closely as possible in the object powerStateDescription.

```
powerStateAllStatesTable(2)
+--powerStateAllStatesEntry(1)
  [powerStateEnergyConsumerId,powerStateName]
  +-- --- SnmpAdminString powerStateName(1)
  +-- r-n Enumeration      powerStateType(2)
  +-- r-n SnmpAdminString powerStateDescription(3)
  +-- r-n Integer32        powerStateAveragePower(4)
  +-- r-n Integer32        powerStateMaximumPower(5)
  +-- r-n TimeTicks        powerStateTotalTime(6)
  +-- r-n TimeStamp        powerStateLastEnterTime(7)
  +-- r-n SnmpAdminString powerStateLastEnterReason(8)
  +-- r-n Counter64        powerStateEnterCount(9)
```

The offered statistics include the total time that the entity spent in a certain power state (`powerStateTotalTime`), the last time at which the entity entered a power state (`powerStateLastEnterTime`), the reason for entering it at the last time (`powerStateLastEnterReason`), the number of times a certain state has been entered (`powerStateEnterCount`), the average power consumed by the entity (`powerStateAveragePower`) and the maximum power consumed by the entity (`powerStateMaximumPower`).

5. Energy MIB

Devices that have instrumentation for measuring electrical energy consumption of entities can implement the Energy MIB module. Entities for which energy consumption is reported can be the entire devices, a component thereof or even an external entity for which the reporting devices observes the energy consumption.

The Energy MIB module defines two tables, the `energyTable` and the `energyPerStateTable`. The first one provides information on the instrumentations and on measured energy consumption of the entity. The second one provides energy consumption information for each individual power state.

5.1. Energy Consumption Table

The first set of managed objects in the `energyTable` are needed to help interpreting the energy consumption readings. These include the power supply type and voltage.

```

energyTable(1)
+--energyEntry(1) [energyConsumerId]
+-- --- Integer32          energyConsumerId(1)
+-- --- ObjectIdentifier   energyConsumerOid(2)
+-- r-n EntitySensorStatus energySensorOperStatus(3)
+-- r-n Unsigned32         energyNominalSupplyVoltage(4)
+-- r-n Enumeration        energyElectricSupplyType(5)
+-- r-n Unsigned32         energyTotalEnergy(6)
+-- r-n UnitMultiplier     energyEnergyUnitMultiplier(7)
+-- r-n Integer32          energyEnergyPrecision(8)
+-- r-n Enumeration        energyMeasurementMethod(9)
+-- r-n TimeStamp          energyDiscontinuityTime(10)
+-- r-n Unsigned32         energySampleInterval(11)
+-- r-n Unsigned32         energyMaxHistory(12)
+-- r-n UnitMultiplier     energyPowerUnitMultiplier(13)
+-- r-n Integer32          energyPowerPrecision(14)
+-- r-n Unsigned32         energyRealPower(15)
+-- r-n Unsigned32         energyPeakRealPower(16)
+-- r-n Unsigned32         energyReactivePower(17)
+-- r-n Unsigned32         energyApparentPower(18)
+-- r-n Integer32          energyPhaseAngle(19)
+-- r-n Integer32          energyPhaseAnglePrecision(20)

```

The main measured values provided by the table are the total energy consumed by the device and the current power (energy consumption rate). For entities supplied with alternating current (AC) there are also objects defined for reporting apparent power, reactive power and phase angle.

Provided energy and power values need to be multiplied by a unit multiplier given by a corresponding unit multiplier object in order to determine a measured value.

Measurements of the total energy consumed by an entity may suffer from interruptions in the continuous measurement of the current energy consumption. In order to indicate such interruptions, object `energyDiscontinuityTime` is provided for indicating the time of the last interruption of total energy measurement.

Time series of energy consumption values for past points in time are stored in the `energyHistoryTable`. Objects `energySampleInterval` and `energyMaxHistory` control the generation of entries in this table, see below.

5.2. Energy Consumption Per Power State Table

The second table in this module is called `energyPerStateTable` and it provides values of total energy consumption per power state in a way

similar to the powerStateTable in the Power State MIB module.

```
energyPerStateTable(2)
+--energyPerStateEntry(1) [energyConsumerId,powerStateName]
+-- r-n Unsigned32 energyPerStateTotalEnergy(1)
```

5.3. Power History Table

The third table in this module is the energyHistoryTable. It stores total energy consumption values for past points in time.

```
energyHistoryTable(3)
+--energyHistoryEntry(1) [energyConsumerId,energyHistoryIndex]
+-- --- Unsigned32 energyHistoryIndex(1)
+-- r-n TimeStamp energyHistoryTimestamp(2)
+-- r-n Unsigned32 energyHistoryTotalEnergy(3)
```

Creation of entries in this table is controlled by the values of corresponding objects energySampleInterval and energyMaxHistory in the energyTable.

Entries are indexed by the the entity (energyConsumerId) and by energyHistoryIndex. The first entry created for a certain entity in the table always has an energyHistoryIndex with a value of 1. Further entries for the same entity get increasing consecutive indices until the maximum index value given by object energyMaxHistory is reached. Then, no further indices will be used, but the entry with the oldest timestamp will be overwritten each time a new entry needs to be created.

A new entry is created with a time difference given by object energySampleInterval after creation of the previous entry. Hence, the difference between timestamps energyHistoryTimestamp of two consecutive entries SHOULD be equal to the value of object energySampleInterval.

6. Battery MIB

Editor's note: The Battery MIB module still uses the entPhysicalIndex from the ENTITY MIB. This will be changed in the next revision.

The third MIB module defined in this document defines objects for reporting information about batteries. The batteryTable contained in the Batter MIB module is again a sparse augment of the Entity MIB module [RFC4133]. It uses one row per battery and require that every battery for which information is provided has its own entry in the entPhysicalTable of the Entity MIB module.

The kind of entity in the entPhysicalTable is indicated by the value of enumeration object entPhysicalClass. Since there is no value called 'battery' defined for this object, it is RECOMMENDED that for batteries the value of this object is chosen to be powerSupply(6).

The batteryTable contains three groups of objects. The first group describes the battery in more detail than the generic objects in the entPhysicalTable. The second group of objects report on the current battery state, if it is charging or discharging, how much it is charged, its remaining capacity, the number of experienced charging cycles, etc.

```
batteryTable(1)
+--batteryEntry(1) [entPhysicalIndex]
   +-- r-n Enumeration batteryType(1)
   +-- r-n Enumeration batteryTechnology(2)
   +-- r-n Unsigned32 batteryNominalVoltage(3)
   +-- r-n Unsigned32 batteryNumberOfCells(4)
   +-- r-n Unsigned32 batteryNominalCapacity(5)
   +-- r-n Unsigned32 batteryRemainingCapacity(6)
   +-- r-n Counter32 batteryChargingCycleCount(7)
   +-- r-n DateAndTime batteryLastChargingCycleTime(8)
   +-- r-n Enumeration batteryState(9)
   +-- r-n Unsigned32 batteryCurrentCharge(10)
   +-- r-n Unsigned32 batteryCurrentChargePercentage(11)
   +-- r-n Unsigned32 batteryCurrentVoltage(12)
   +-- r-n Integer32 batteryCurrentCurrent(13)
   +-- r-n Unsigned32 batteryLowAlarmPercentage(14)
   +-- r-n Unsigned32 batteryLowAlarmVoltage(15)
   +-- r-n Unsigned32 batteryReplacementAlarmCapacity(16)
   +-- r-n Unsigned32 batteryReplacementAlarmCycles(17)
```

The third group of objects in this table indicates thresholds which can be used to raise an alarm if a property of the battery exceeds one of them. Raising an alarm may include sending a notification. The Battery MIB defines two notifications, one indicating a low battery charging state and one indicating an aged battery that may need to be replaced.

7. Relationship to Other MIB Modules

The three MIB modules described above relate to a number of existing standard MIB modules and complements them where necessary.

This section needs to be revised.

8. Definitions

8.1. Power State MIB

```
POWER-STATE-MIB DEFINITIONS ::= BEGIN
```

```
IMPORTS
```

```
    MODULE-IDENTITY, OBJECT-TYPE, NOTIFICATION-TYPE,  
    mib-2, Integer32, Counter64, TimeTicks  
        FROM SNMPv2-SMI -- RFC2578  
    TimeStamp  
        FROM SNMPv2-TC -- RFC2579  
    MODULE-COMPLIANCE, OBJECT-GROUP, NOTIFICATION-GROUP  
        FROM SNMPv2-CONF -- RFC2580  
    SnmpAdminString  
        FROM SNMP-FRAMEWORK-MIB; -- RFC3411
```

```
powerStateMIB MODULE-IDENTITY
```

```
    LAST-UPDATED "201010231200Z" -- 23 October 2010  
    ORGANIZATION "IETF OPSAWG Working Group"  
    CONTACT-INFO  
        "General Discussion: opsawg@ietf.org  
        To Subscribe: https://www.ietf.org/mailman/listinfo/opsawg  
        Archive: http://www.ietf.org/mail-archive/web/opsawg
```

```
    Co-editor:
```

```
        Juergen Quittek  
        NEC Europe Ltd.  
        NEC Laboratories Europe  
        Kurfuersten-Anlage 36  
        69115 Heidelberg  
        Germany  
        Tel: +49 6221 4342-115  
        Email: quittek@neclab.eu
```

```
    Co-editor:
```

```
        Thomas Dietz  
        NEC Europe Ltd.  
        NEC Laboratories Europe  
        Kurfuersten-Anlage 36  
        69115 Heidelberg  
        Germany  
        Phone: +49 6221 4342-128  
        Email: Thomas.Dietz@neclab.eu"
```

```
DESCRIPTION
```

```
    "This MIB module defines a set of objects for monitoring  
    the power state of managed entitites."
```

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

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this MIB module is part of RFC yyyy; see the RFC itself for full legal notices."

-- replace yyyy with actual RFC number & remove this notice

-- Revision history

REVISION "201010231200Z" -- 23 October 2010
DESCRIPTION

"Initial version, published as RFC yyyy."

-- replace yyyy with actual RFC number & remove this notice

::= { mib-2 9991 }

-- xxx to be assigned by IANA.

-- *****
-- Top Level Structure of the MIB module
-- *****

powerStateNotifications OBJECT IDENTIFIER ::= { powerStateMIB 0 }
powerStateObjects OBJECT IDENTIFIER ::= { powerStateMIB 1 }
powerStateConformance OBJECT IDENTIFIER ::= { powerStateMIB 2 }

=====
-- 1. Object Definitions
=====

-- 1.1. Actual Power State Table

powerStateTable OBJECT-TYPE
SYNTAX SEQUENCE OF PowerStateEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
"This table provides information on the current power state of managed entities.

The table is indexed by an ID of the entity on which power

state information is provided. IDs can be provided by another MIB module, such as the ENERGY AWARE MIB module or the ENTITY MIB module. If not ID provisioning from other MIB modules is available, the table can only have one entry for reporting the local power state of the device that tuns an instance of this table."

```
::= { powerStateObjects 1 }
```

```
powerStateEntry OBJECT-TYPE
```

```
SYNTAX          PowerStateEntry
```

```
MAX-ACCESS      not-accessible
```

```
STATUS          current
```

```
DESCRIPTION
```

"An entry providing information on the current power state of an entity."

```
INDEX { powerStateEnergyConsumerId }
```

```
::= { powerStateTable 1 }
```

```
PowerStateEntry ::=
```

```
SEQUENCE {
```

```
    powerStateEnergyConsumerId      Integer32,
```

```
    powerStateEnergyConsumerOid     OBJECT IDENTIFIER,
```

```
    powerStateOperationalState      SnmpAdminString,
```

```
    powerStateAdminState            SnmpAdminString
```

```
}
```

```
powerStateEnergyConsumerId OBJECT-TYPE
```

```
SYNTAX          Integer32 (0..2147483647)
```

```
MAX-ACCESS      not-accessible
```

```
STATUS          current
```

```
DESCRIPTION
```

"An integer that identifies an entity that is subject of power state monitoring. Index values MUST be locally unique for each identified entity."

If an implementation of the ENERGY AWARE MIB module is available in the local SNMP context, then the same index of an entity MUST be chosen as assigned to the entity by object pmIndex in the ENERGY AWARE MIB module. In this case, entities without an assigned value for pmIndex cannot be indexed by the powerCurrentStateTable.

If there is no implementation of the ENERGY AWARE MIB module but one of the ENTITY MIB module is available in the local SNMP context, then the same index of an entity MUST be chosen as assigned to the entity by object entPhysicalIndex in the ENTITY MIB module. In this case, entities without an assigned value for pmIndex cannot be indexed by the

powerCurrentStateTable.

If neither the ENERGY AWARE MIB module nor of the ENTITY MIB module is available in the local SNMP context, then this MIB module may choose identity values from a further MIB module providing entity identities. In this case the value for each pmIndex must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.

In case that no other MIB module has been chosen for providing entity identities, power state can be reported exclusively for the local device on which this table is instantiated. Then this table will have a single entry only and an index value of 0 MUST be used.

The identity provisioning method that has been chosen can be retrieved by reading the value of powerStateEnergyConsumerOid. In case of identities provided by the ENERGY AWARE MIB module, this OID points to an existing instance of pmIndex, in case of the ENTITY MIB, the object points to a valid instance of entPhysicalIndex, and in a similar way, it points to a value of another MIB module if this is used for identifying entities. If no other MIB module has been chosen for providing entity identities, then the value of powerStateEnergyConsumerOid MUST be 0.0 (zeroDotZero)."

::= { powerStateEntry 1 }

powerStateEnergyConsumerOid OBJECT-TYPE

SYNTAX OBJECT IDENTIFIER

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"An OID that identifies an entity that is subject of power state monitoring. The value MUST be an OID that points to an existing managed object or 0.0 (zeroDotZero)."

If another MIB module is chosen for providing identities for managed entities, then the value of this object points to an existing instance of an entity identifier, such as an instance of pmIndex in the ENERGY AWARE MIB or an instance of entPhysicalIndex in the ENTITY MIB module.

If power state information is provided only for the local device on which this table is instantiated, then the value of this object MUST be 0.0 (zeroDotZero)."

::= { powerStateEntry 2 }

```
powerStateOperationalState OBJECT-TYPE
    SYNTAX      SnmpAdminString (SIZE(1..32))
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "This object indicates the current power state of the
        entity. The given SnmpAdminString MUST match the
        powerStateName object of an entry in the
        powerStateAllStatesTable."
    ::= { powerStateEntry 3 }

powerStateAdminState OBJECT-TYPE
    SYNTAX      SnmpAdminString (SIZE(0..32))
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "This object indicates the desired power state of the
        entity. This object may be set by a network management
        system in order to request changing the actual power state
        to the desired one.

        If this object has not been set by an administrative action
        requesting a certain power state, then its value is an
        empty string of length 0."
    ::= { powerStateEntry 4 }

-----
-- 1.2. All Power States Table
-----

powerStateAllStatesTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF PowerStateAllStatesEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "This table provides information on all available power
        states of managed entities.

        The table extends the powerStateTable by sharing the first
        index. The first index serves for identifying an entity for
        which power state information is provided. The second index
        identifies a single power state by its name."
    ::= { powerStateObjects 2 }

powerStateAllStatesEntry OBJECT-TYPE
    SYNTAX      PowerStateAllStatesEntry
    MAX-ACCESS   not-accessible
    STATUS       current
```

DESCRIPTION

"Power state information about this physical entity."

INDEX { powerStateEnergyConsumerId, powerStateName }
 ::= { powerStateAllStatesTable 1 }

PowerStateAllStatesEntry ::=

```
SEQUENCE {
    powerStateName          SnmpAdminString,
    powerStateType          INTEGER,
    powerStateDescription   SnmpAdminString,
    powerStateAveragePower  Integer32,
    powerStateMaximumPower  Integer32,
    powerStateTotalTime     TimeTicks,
    powerStateLastEnterTime TimeStamp,
    powerStateLastEnterReason SnmpAdminString,
    powerStateEnterCount    Counter64
}
```

powerStateName OBJECT-TYPE

SYNTAX SnmpAdminString (SIZE(1..32))

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"This index should only be created for power states that are actually implemented by the entity that is identified by the first index powerStateEnergyConsumerOid.

This index is the name of the power state and is limited to 32 characters.

If possible the name SHOULD already give a rough idea of the characteristic of this power state."

::= { powerStateAllStatesEntry 1 }

powerStateType OBJECT-TYPE

```
SYNTAX INTEGER {
    unknown(0),
    off(1),
    nonOperational(2),
    operational(3)
}
```

 -- Open issue: Shall we replace the syntax by textual convention
 -- PowerMonitorLevel from draft-claise-energy-monitoring-mib?

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"Object classifies the power state. It helps to clearly distinguish non-operational power states (sleep, standby, etc.) from operational ones. In a nonOperational(2) state an entity provides non of its primary services except for bringing it into operational(3) states or off(1) states.

A device in state off(1) cannot report its state on its own. But state off(1) may be reported by managed devices reporting on the power state of other managed devices."

::= { powerStateAllStatesEntry 2 }

powerStateDescription OBJECT-TYPE

SYNTAX SnmpAdminString

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"Power states are identified by their names. However, semantics of power states may vary between different entities. Reasons for variations can be different hardware and software architectures of managed devices.

Object powerStateDescription SHOULD describe the power state and its characteristics as closely as possible."

::= { powerStateAllStatesEntry 3 }

powerStateAveragePower OBJECT-TYPE

SYNTAX Integer32

UNITS "milliwatt"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the average power (energy consumption rate) in milliwatt at the electrical power supply of the entity in the power state indicated by powerStateName.

A value of -1 indicates that the average power in this state is unknown."

::= { powerStateAllStatesEntry 4 }

powerStateMaximumPower OBJECT-TYPE

SYNTAX Integer32

UNITS "milliwatt"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the maximum power (energy consumption rate) in milliwatt at the electrical power supply of the

entity in the power state indicated by powerStateName.

A value of -1 indicates that the maximum power in this state is unknown."

::= { powerStateAllStatesEntry 5 }

powerStateTotalTime OBJECT-TYPE

SYNTAX TimeTicks

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the total time in hundreds of seconds that the entity has been in the state indicated by index powerStateName."

::= { powerStateAllStatesEntry 6 }

-- Open issue: Shall we use DateAndTime instead of timeTicks?

powerStateLastEnterTime OBJECT-TYPE

SYNTAX TimeStamp

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This time stamp object indicates the last time a which the entity entered the state indicated by index powerStateName."

::= { powerStateAllStatesEntry 7 }

powerStateLastEnterReason OBJECT-TYPE

SYNTAX SnmpAdminString

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This string object describes the reason for the last power state transition into the power state indicated by index powerStateName."

::= { powerStateAllStatesEntry 8 }

powerStateEnterCount OBJECT-TYPE

SYNTAX Counter64

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates how often the entity indicated by index entPhysicalIndex entered the power state indicated by index powerStateName."

::= { powerStateAllStatesEntry 9 }

```
=====
-- 2. Notifications
=====

powerStateChangeEvent NOTIFICATION-TYPE
    OBJECTS      { powerStateLastEnterReason }
    STATUS       current
    DESCRIPTION
        "This notification can be generated when the power state of
        an entity changes.

        Note that the state that has been entered is indicated by
        the OID of object powerStateLastEnterReason."
    ::= { powerStateNotifications 1 }

=====
-- 3. Conformance Information
=====

powerStateCompliances OBJECT IDENTIFIER
    ::= { powerStateConformance 1 }
powerStateGroups      OBJECT IDENTIFIER
    ::= { powerStateConformance 2 }

-----
-- 3.1. Compliance Statements
-----

powerCompliance MODULE-COMPLIANCE
    STATUS       current
    DESCRIPTION
        "The compliance statement for implementations of the
        POWER-STATE-MIB module.

        A compliant implementation MUST implement the objects
        defined in the mandatory group powerRequiredGroup."
    MODULE      -- this module
    MANDATORY-GROUPS { powerStateRequiredGroup }
    GROUP       powerStateNotificationsGroup
    DESCRIPTION
        "A compliant implementation does not have to implement
        the powerNotificationsGroup."
    ::= { powerStateCompliances 1 }

-----
-- 3.2. MIB Grouping
-----
```

```
powerStateRequiredGroup OBJECT-GROUP
    OBJECTS {
        powerStateOperationalState,
        powerStateAdminState,
        powerStateType,
        powerStateDescription,
        powerStateTotalTime,
        powerStateLastEnterTime,
        powerStateLastEnterReason,
        powerStateEnterCount,
        powerStateAveragePower,
        powerStateMaximumPower
    }
    STATUS      current
    DESCRIPTION
        "A compliant implementation MUST implement the objects
        contained in this group."
    ::= { powerStateGroups 1 }

powerStateNotificationsGroup NOTIFICATION-GROUP
    NOTIFICATIONS { powerStateChangeEvent }
    STATUS      current
    DESCRIPTION
        "A compliant implementation does not have to implement the
        notification contained in this group."
    ::= { powerStateGroups 2 }
END
```

8.2. Energy MIB

```
ENERGY-MIB DEFINITIONS ::= BEGIN

IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE, mib-2,
    Unsigned32, Integer32
        FROM SNMPv2-SMI
        -- RFC2578
    TimeStamp
        FROM SNMPv2-TC
        -- RFC2579
    MODULE-COMPLIANCE, OBJECT-GROUP
        FROM SNMPv2-CONF
        -- RFC2580
    EntitySensorStatus
        FROM ENTITY-SENSOR-MIB
        -- RFC3433
    powerStateName
        FROM POWER-STATE-MIB
    UnitMultiplier
        FROM POWER-MONITOR-MIB;

energyMIB MODULE-IDENTITY
```


LAST-UPDATED "201010231200Z" -- 23 October 2010
ORGANIZATION "IETF OPSAWG Working Group"
CONTACT-INFO
"General Discussion: opsawg@ietf.org
To Subscribe: <https://www.ietf.org/mailman/listinfo/opsawg>
Archive: <http://www.ietf.org/mail-archive/web/opsawg>

Co-editor:

Juergen Quittek
NEC Europe Ltd.
NEC Laboratories Europe
Kurfuersten-Anlage 36
69115 Heidelberg
Germany
Tel: +49 6221 4342-115
Email: quittek@neclab.eu

Co-editor:

Thomas Dietz
NEC Europe Ltd.
NEC Laboratories Europe
Kurfuersten-Anlage 36
69115 Heidelberg
Germany
Phone: +49 6221 4342-128
Email: Thomas.Dietz@neclab.eu

DESCRIPTION

"This MIB module defines a set of objects for monitoring the energy consumption of networked devices and their components.

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

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this MIB module is part of RFC yyyy; see the RFC itself for full legal notices."

-- replace yyyy with actual RFC number & remove this notice

-- Revision history

```

    REVISION      "201010231200Z"          -- 23 October 2010
    DESCRIPTION
        "Initial version, published as RFC yyyy."
-- replace yyyy with actual RFC number & remove this notice

    ::= { mib-2 9992 }
-- yyy to be assigned by IANA.

--*****
-- Top Level Structure of the MIB module
--*****

energyObjects      OBJECT IDENTIFIER ::= { energyMIB 1 }
energyConformance  OBJECT IDENTIFIER ::= { energyMIB 2 }

=====
-- 1. Object Definitions
=====

-----
-- 1.1. Energy Consumption Table
-----

energyTable  OBJECT-TYPE
    SYNTAX      SEQUENCE OF EnergyEntry
    MAX-ACCESS   not-accessible
    STATUS      current
    DESCRIPTION
        "This table provides information on the current and
        accumulated energy consumption of entities.

        The table is indexed by an ID of the entity on which
        energy information is provided. IDs can be provided by
        another MIB module, such as the ENERGY AWARE MIB module
        or the ENTITY MIB module. If not ID provisioning from
        other MIB modules is available, the table can only have
        one entry for reporting the local power state of the
        device that runs an instance of this table."
    ::= { energyObjects 1 }

energyEntry  OBJECT-TYPE
    SYNTAX      EnergyEntry
    MAX-ACCESS   not-accessible
    STATUS      current
    DESCRIPTION
        "An entry providing information on the energy consumption
        of a physical entity."
    INDEX       { energyConsumerId }
    ::= { energyTable 1 }

```

```

EnergyEntry ::=
  SEQUENCE {
    energyConsumerId          Integer32,
    energyConsumerOid         OBJECT IDENTIFIER,
    energySensorOperStatus    EntitySensorStatus,
    energyNominalSupplyVoltage Unsigned32,
    energyElectricSupplyType  INTEGER,
    energyTotalEnergy         Unsigned32,
    energyEnergyUnitMultiplier UnitMultiplier,
    energyEnergyPrecision     Integer32,
    energyMeasurementMethod   INTEGER,
    energyDiscontinuityTime   TimeStamp,
    energySampleInterval      Unsigned32,
    energyMaxHistory          Unsigned32,
    energyPowerUnitMultiplier UnitMultiplier,
    energyPowerPrecision      Integer32,
    energyRealPower           Unsigned32,
    energyPeakRealPower       Unsigned32,
    energyReactivePower        Unsigned32,
    energyApparentPower        Unsigned32,
    energyPhaseAngle          Integer32,
    energyPhaseAnglePrecision Integer32
  }

```

energyConsumerId OBJECT-TYPE

SYNTAX Integer32 (0..2147483647)

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"An integer that identifies an entity that is subject of energy monitoring. Index values MUST be locally unique for each identified entity.

If an implementation of the ENERGY AWARE MIB module is available in the local SNMP context, then the same index of an entity MUST be chosen as assigned to the entity by object pmIndex in the ENERGY AWARE MIB module. In this case, entities without an assigned value for pmIndex cannot be indexed by the powerCurrentStateTable.

If there is no implementation of the ENERGY AWARE MIB module but one of the ENTITY MIB module is available in the local SNMP context, then the same index of an entity MUST be chosen as assigned to the entity by object entPhysicalIndex in the ENTITY MIB module. In this case, entities without an assigned value for pmIndex cannot be indexed by the powerCurrentStateTable.

If neither the ENERGY AWARE MIB module nor of the ENTITY MIB module is available in the local SNMP context, then this MIB module may choose identity values from a further MIB module providing entity identities. In this case the value for each pmIndex must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.

In case that no other MIB module has been chosen for providing entity identities, power state can be reported exclusively for the local device on which this table is instantiated. Then this table will have a single entry only and an index value of 0 MUST be used.

The identity provisioning method that has been chosen can be retrieved by reading the value of object powerStateEnergyConsumerOid. In case of identities provided by the ENERGY AWARE MIB module, this OID points to an existing instance of pmIndex, in case of the ENTITY MIB, the object points to a valid instance of entPhysicalIndex, and in a similar way, it points to a value of another MIB module if this is used for identifying entities. If no other MIB module has been chosen for providing entity identities, then the value of powerStateEnergyConsumerOid MUST be 0.0 (zeroDotZero)."

```
::= { energyEntry 1 }
```

energyConsumerOid OBJECT-TYPE

SYNTAX OBJECT IDENTIFIER

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"An OID that identifies an entity that is subject of energy monitoring. The value MUST be an OID that points to an existing managed object or 0.0 (zeroDotZero)."

If another MIB module is chosen for providing identities for managed entities, then the value of this object points to an existing instance of an entity identifier, such as an instance of pmIndex in the ENERGY AWARE MIB or an instance of entPhysicalIndex in the ENTITY MIB module.

If power state information is provided only for the local device on which this table is instantiated, then the value of this object MUST be 0.0 (zeroDotZero)."

```
::= { energyEntry 2 }
```

energySensorOperStatus OBJECT-TYPE

```
SYNTAX      EntitySensorStatus
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "This object provides the operational status of the
    sensor that is used for measuring the energy consumption
    of the entity indicated by energyConsumerId."
 ::= { energyEntry 3 }
```

energyNominalSupplyVoltage OBJECT-TYPE

```
SYNTAX      Unsigned32
UNITS        "millivolt"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "This object provides the nominal voltage of the power
    supply of the entity. It is provided in units of
    millivolt (mV).

    The nominal voltage actual of an entity is assumed to be
    fixed, while the actual power supply voltage may vary over
    time, for example, caused by changing load conditions.

    A value of 0 indicates that the nominal supply voltage
    is unknown."
 ::= { energyEntry 4 }
```

energyElectricSupplyType OBJECT-TYPE

```
SYNTAX      INTEGER {
                    alternatingCurrent(1),
                    directCurrent(2),
                    unknown(3)
                }
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "This object indicates the type of electrical power
    supply for the entity. It is used for distinguishing
    between alternating current (AC) supply and direct
    current (DC) supply."
 ::= { energyEntry 5 }
```

energyTotalEnergy OBJECT-TYPE

```
SYNTAX      Unsigned32
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "This object indicates the total consumed energy measured
```

at the electrical power supply of the entity.

In order to determine the measured value in watt hours, the value of this object needs to be multiplied by a unit multiplier given by the value of object `energyEnergyUnitMultiplier`.

Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of `energyDiscontinuityTime`."

::= { energyEntry 6 }

`energyEnergyUnitMultiplier` OBJECT-TYPE

SYNTAX UnitMultiplier

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object provides unit multiplier for measured energy values. Reported values need to be multiplied with this multiplier in order to determine the measured value in watt hours.

This object serves as unit multiplier for objects `energyTotalEnergy`, `energyPSTotalEnergy`,

..."

::= { energyEntry 7 }

`energyEnergyPrecision` OBJECT-TYPE

SYNTAX Integer32 (0..10000)

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates a the precision of a measured energy value. The precision is indicated as a percentage value, in 100ths of a percent. A value of 0 indicates that the precision is unknown or not applicable to the measured value.

This object serves precision indicator for the values provided by objects `energyTotalEnergy`, `energyPSTotalEnergy`, ..."

::= { energyEntry 8 }

`energyMeasurementMethod` OBJECT-TYPE

SYNTAX INTEGER {
 directEnergyMeasurement(1),
 powerOversampling(2),

```

        powerSampling(3),
        loadBasedEstimation(4),
        deviceBasedEstimation(5),
        unknown(6)
    }
MAX-ACCESS    read-only
STATUS        current
DESCRIPTION
    "This object indicates the method used for measuring energy
    consumption. A device may not be equipped with capabilities
    to measure its energy consumption directly, but rather
    relies on other input in order to conduct more or less
    precise estimations of its power consumption.

    The measurement methods concerns values of objects
    energyTotalEnergy, energyPSTotalEnergy, and
    energyPowerHistoryAverageValue.

    Five different measurement methods are specified.

    - directEnergyMeasurement(1) indicates that the entity is
      instrumented to directly measure its energy consumption.

    - powerOversampling(2) indicates that energy is measured
      by sampling power values more frequently than indicated
      by the value of object energySampleInterval.

    - powerSampling(3) indicates that energy is measured
      by sampling power values according to the value of object
      energySampleInterval.

    - loadBasedEstimation(4) indicates that power is estimated
      based on measurements of the load of the entity.

    - deviceBasedEstimation(5) indicates that power is estimated
      based on static properties of the entity. In this case,
      reported power only depends on the power state of the
      devices as indicated by object powerCurrentState in the
      powerCurrentStateTable of the Power State MIB module."
 ::= { energyEntry 9 }

```

energyDiscontinuityTime OBJECT-TYPE

```

SYNTAX        TimeStamp
MAX-ACCESS    read-only
STATUS        current
DESCRIPTION

```

"The value of sysUpTime on the most recent occasion at which
any one or more of this entity's energy consumption counters

suffered a discontinuity. The relevant counters are energyTotalEnergy and energyPerStateTotalEnergy. If no such discontinuities have occurred since the last re-initialization of the local management subsystem, then this object contains a zero value."

```
::= { energyEntry 10 }
```

energySampleInterval OBJECT-TYPE

SYNTAX Unsigned32
UNITS "milliseconds"
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object indicates is the difference of time stamps between two consecutive entries in the energyHistoryTable for this entity.

The interval lenght provided by this object indicates the or maximum interval length (or minimal sampling rate) at which the power sensor measures values of the current power. Implementations of the Energy MIB module may choose higher sampling rates (or shorter sampling intervals) in order to provide higher precision of the measurement. Preferably, shorter intervals may be chosen such that the sampling interval indicated by this object is a multiple of the actual sampling interval.

The sampling interval is provided in units of microseconds.

A value of 0 indicates that the sampling interval applied by the sensor is unknown or not constant."

```
::= { energyEntry 11 }
```

energyMaxHistory OBJECT-TYPE

SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object indicates is the maximum number of corresponding entries in the energyPowerHistoryTable. An entry in the energyHistoryTable is corresponding if it has the same value for object energyConsumerId as index.

An implementation of the Energy MIB module will remove the oldest correaponding entry in the energyHistoryTable to allow the addition of a new entry once the number of corresponding entries in the energyHistoryTable

reaches this value.

Entries are added to the energyHistoryTable until energyMaxHistory is reached before entries begin to be removed.

A value of 0 for this object disables creation of corresponding energyHistoryTable entries."

DEFVAL { 0 }
 ::= { energyEntry 12 }

energyPowerUnitMultiplier OBJECT-TYPE

SYNTAX UnitMultiplier

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object provides unit multiplier for measured energy values. Reported values need to be multiplied with this multiplier in order to determine the measured value in watt hours.

This object serves as unit multiplier for the values provided by objects energyRealPower, energyPeakRealPower, energyReactivePower, and energyApparentPower."

::= { energyEntry 13 }

energyPowerPrecision OBJECT-TYPE

SYNTAX Integer32 (0..10000)

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates a the precision of a measured power value. The precision is indicated as a percentage value, in 100ths of a percent. A value of 0 indicates that the precision is unknown or not applicable to the measured value.

This object serves precision indicator for the values provided by objects energyRealPower, energyPeakRealPower, energyReactivePower, and energyApparentPower."

::= { energyEntry 14 }

energyRealPower OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the current real power value

at the electrical supply of the entity indicated by index energyConsumerId.

In order to determine the measured value in watts, the value of this object needs to be multiplied by a unit multiplier given by the value of object energyEnergyUnitMultiplier.

Measured values of this object are stored in the energyPowerTable with a rate determined by object energySampleInterval."

::= { energyEntry 15 }

energyPeakRealPower OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the highest observed value for object energyRealPower since the last re-initialization of the management system.

In order to determine the measured value in watts, the value of this object needs to be multiplied by a unit multiplier given by the value of object energyEnergyUnitMultiplier."

::= { energyEntry 16 }

energyReactivePower OBJECT-TYPE

SYNTAX Unsigned32

UNITS "volt-amperes reactive"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the current reactive power value at the electrical supply of the entity indicated by index energyConsumerId.

In order to determine the measured value in volt-amperes (var), the value of this object needs to be multiplied by a unit multiplier given by the value of object energyEnergyUnitMultiplier.

The value provided by this object is only useful if the value of object energySupplyType is alternatingCurrent(1). In this case it is RECOMMENDED that at least one of the three values energyReactivePower, energyApparentPowerScale, and energyPhaseAngle

are provided.

If object `energyElectricSupplyType` of this row has a value other than `alternatingCurrent(1)`, then the value of this object MUST be 0.

If object `energyElectricSupplyType` of this row has the value `alternatingCurrent(1)` and if no value for the current reactive power is provided, then the value of this object MUST be 0xFFFF."

::= { energyEntry 17 }

`energyApparentPower` OBJECT-TYPE

SYNTAX Unsigned32

UNITS "volt-amperes"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the current apparent power value measured in volt-ampere (VA) at the electrical supply of the entity for a time interval indicated by object `energySampleInterval`.

The value provided by this object is only useful if the value of object `energySupplyType` is `alternatingCurrent(1)`. In this case it is RECOMMENDED that at least one of the three values `energyReactivePower`, `energyApparentPowerScale`, and `energyPhaseAngle` are provided.

Scale and precision of the value are indicated by objects `energyPowerScale` and `energyPowerPrecision`.

If object `energyElectricSupplyType` of this row has a value other than `alternatingCurrent(1)`, then the value of this object MUST be equal to the value of object `energyRealPower`.

If object `energyElectricSupplyType` of this row has the value `alternatingCurrent(1)` and if no value for the current apparent power is provided, then the value of this object MUST be -10000000000."

::= { energyEntry 18 }

`energyPhaseAngle` OBJECT-TYPE

SYNTAX Integer32 (-1..360000)

UNITS "millidegrees"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the current phase angle value measured at the electrical supply of the entity for a time interval indicated by object energySampleInterval.

The value provided by this object is only useful if the value of object energySupplyType is alternatingCurrent(1). In this case it is RECOMMENDED that at least one of the three values energyReactivePower, energyApparentPowerScale, and energyPhaseAngle are provided.

The value is provided in units of millidegree (one thousands of a degree. This is equivalent to an associated object of type EntitySensorDataScale with the value of milli(8) and an associated object of type EntitySensorPrecision with a value of 0.

The minimum value for this object when indicating an actual angle is 0, the maximum value is 360000.

The maximum error of of the value is indicated by object energyPhaseAngleMaxError.

If object energyElectricSupplyType of this row has a value other than alternatingCurrent(1), then the value of this object MUST be 0.

If object energyElectricSupplyType of this row has the value alternatingCurrent(1) and if no value for the phase angle is provided, then the value of this object MUST be -1."

::= { energyEntry 19 }

energyPhaseAnglePrecision OBJECT-TYPE

SYNTAX Integer32 (0..10000)

UNITS "millidegrees"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates a the precision of a measured phase angle value. The precision is indicated as a percentage value, in 100ths of a percent. A value of 0 indicates that the precision is unknown or not applicable to the measured value.

This object serves precision indicator for the values

```
        provided by object energyPhaseAngle."
 ::= { energyEntry 20 }
```

```
-----
-- 1.2. Energy Consumption Per Power State Table
-----
```

```
energyPerStateTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF EnergyPerStateEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table provides information on the accumulated energy
        consumption of an entity.

        This table extends the energyTable by sharing the
        first index. The first index serves for identifying an
        entity for which energy information is provided. The second
        index identifies a single power state by its name."
    ::= { energyObjects 2 }
```

```
energyPerStateEntry OBJECT-TYPE
    SYNTAX      EnergyPerStateEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "Energy consumption information per power state for a
        physical entity."
    INDEX { energyConsumerId, powerStateName }
    ::= { energyPerStateTable 1 }
```

```
EnergyPerStateEntry ::=
    SEQUENCE {
        energyPerStateTotalEnergy      Unsigned32
    }
```

```
energyPerStateTotalEnergy OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object indicates the total consumed energy value
        at the electrical supply of the entity indicated by index
        energyConsumerId while being in a specific power state
        indicated by index powerStateName.

        In order to determine the measured value in watts, the value
        of this object needs to be multiplied by a unit multiplier
```

given by the value of object
energyEnergyUnitMultiplier of table
energyTable.

Discontinuities in the value of this counter can occur at
re-initialization of the management system, and at other
times as indicated by the value of
energyDiscontinuityTime."
 ::= { energyPerStateEntry 1 }

-- 1.3. Energy Power History Table

energyHistoryTable OBJECT-TYPE

SYNTAX SEQUENCE OF EnergyHistoryEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"This table stores results of energy consumption
measurements for multiple entities.

This table extends the energyTable by sharing the
first index. The first index serves for identifying an
entity for which energy information is provided. The second
index energyHistoryIndex identifies a single measurement
consisting of an energy consumption value and a timestamp.

Creation of entries in this row is controlled individually
for each entity by two parameters: energyMaxHistory and
energySamplingInterval.

The energySamplingInterval controls the difference in time
between the creation of two consecutive entries in this
table. Object energyMaxHistory limits the number of entries
in this table that can be created for the corresponding
entity.

An implementation of the Energy MIB module will remove the
oldest entry for an entity in the energyHistoryTable to
allow the addition of a new entry once the number of
entries for this entity reaches the value indicated by
object energyMaxHistory.

Entries for a specific entity are added to this table
until energyMaxHistory is reached before
entries begin to be removed.

Entries for the same entity are indexed by energyHistoryIndex. The first entry for an entity MUST have an index value of 1. Further new entries MUST be indexed by consecutive numbers in the order in which they are created until the value of energyMaxHistory is reached. Then no further new indices will be assigned, but existing ones will be re-used."

::= { energyObjects 3 }

energyHistoryEntry OBJECT-TYPE

SYNTAX EnergyHistoryEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"An entry indicating consumed energy for an entity at a certain point in time."

INDEX { energyConsumerId, energyHistoryIndex }

::= { energyHistoryTable 1 }

EnergyHistoryEntry ::=

SEQUENCE {

energyHistoryIndex Unsigned32,

energyHistoryTimestamp TimeStamp,

energyHistoryTotalEnergy Unsigned32

}

energyHistoryIndex OBJECT-TYPE

SYNTAX Unsigned32 (1..4294967295)

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"The index for this entry per entity.

Values of this index MUST be unique per entity used as first index."

::= { energyHistoryEntry 1 }

energyHistoryTimestamp OBJECT-TYPE

SYNTAX TimeStamp

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the time at which the energy consumption value provided by object energyHistoryTotalEnergy was measured."

::= { energyHistoryEntry 2 }

energyHistoryTotalEnergy OBJECT-TYPE

SYNTAX Unsigned32

```
MAX-ACCESS    read-only
STATUS        current
DESCRIPTION
    "This object indicates the total consumed energy measured
    at the electrical power supply of the entity.

    In order to determine the measured value in watt hours,
    the value of this object needs to be multiplied by a unit
    multiplier given by the value of object
    energyEnergyUnitMultiplier in the corresponding entry
    for this entity in table energyTable.

    Discontinuities in the value of this counter can occur at
    re-initialization of the management system, and at other
    times as indicated by the value of
    energyDiscontinuityTime in the corresponding entry
    for this entity in table energyTable."
 ::= { energyHistoryEntry 3 }
```

```
=====
-- 2. Conformance Information
=====
```

```
energyCompliances OBJECT IDENTIFIER ::= { energyConformance 1 }
energyGroups       OBJECT IDENTIFIER ::= { energyConformance 2 }
```

```
-----
-- 2.1. Compliance Statements
-----
```

```
energyCompliance MODULE-COMPLIANCE
STATUS        current
DESCRIPTION
    "The compliance statement for implementations of the
    ENERGY-MIB module.

    A compliant implementation MUST implement the objects
    defined in the mandatory group energyRequiredGroup.

    If one of the entities for which energy consumption is
    reported are supplied by alternating current (AC) then it
    is recommended that not just real power is reported
    (REQUIRED) but it is also RECOMMENDED that at least one
    of three other related values (reactive power, apparent
    power, and phase angle) is reported by implementing at least
    one of the three groups energyReactivePowerGroup,
    energyApparentPowerGroup, and energyPhaseAngleGroup."
```



```
MODULE -- this module
MANDATORY-GROUPS { energyRequiredGroup }

GROUP energyPowerHistoryGroup
DESCRIPTION
    "This group is only needed for implementations that
    support storing time series of measured power values
    in the energyPowerHistoryTable."

GROUP energyACGroup
DESCRIPTION
    "This group is only needed for implementations that report
    consumption of electric energy provided by alternating
    current (AC) supply.

    Implementations for devices supplied with direct current (DC)
    only and implementations that do only report real power
    reporting for alternative current do not need to implement
    objects in this group."

GROUP energyReactivePowerGroup
DESCRIPTION
    "Information provided by elements in this group is redundant
    to information provided by elements in the
    energyApparentPowerGroup and the energyPhaseAngleGroup.

    For compliant implementations that report consumption of
    electric energy provided by alternating current (AC) supply
    it is RECOMMENDED to at least one of the three groups
    energyReactivePowerGroup, energyApparentPowerGroup, and
    energyPhaseAngleGroup."

GROUP energyApparentPowerGroup
DESCRIPTION
    "Information provided by elements in this group is redundant
    to information provided by elements in the
    energyReactivePowerGroup and the energyPhaseAngleGroup.

    For compliant implementations that report consumption of
    electric energy provided by alternating current (AC) supply
    it is RECOMMENDED to at least one of the three groups
    energyReactivePowerGroup, energyApparentPowerGroup, and
    energyPhaseAngleGroup."

GROUP energyPhaseAngleGroup
DESCRIPTION
    "Information provided by elements in this group is redundant
    to information provided by elements in the
```

energyReactivePowerGroup and the energyApparentPowerGroup.

For compliant implementations that report consumption of electric energy provided by alternating current (AC) supply it is RECOMMENDED to at least one of the three groups energyReactivePowerGroup, energyApparentPowerGroup, and energyPhaseAngleGroup."

::= { energyCompliances 1 }

-- 2.2. Object Grouping

energyRequiredGroup OBJECT-GROUP

OBJECTS {
 energySensorOperStatus,
 energyNominalSupplyVoltage,
 energyElectricSupplyType,
 energyTotalEnergy,
 energyEnergyUnitMultiplier,
 energyEnergyPrecision,
 energyMeasurementMethod,
 energyDiscontinuityTime,
 energyPowerUnitMultiplier,
 energyPowerPrecision,
 energyRealPower,
 energyPeakRealPower,
 energyPerStateTotalEnergy
}

STATUS current

DESCRIPTION

"A compliant implementation MUST implement the objects contained in this group."

::= { energyGroups 1 }

energyPowerHistoryGroup OBJECT-GROUP

OBJECTS {
 energySampleInterval,
 energyMaxHistory,
 energyHistoryTimestamp,
 energyHistoryTotalEnergy
}

STATUS current

DESCRIPTION

"The group of object for reporting details of AC power measurement."

::= { energyGroups 2 }

```
energyACGroup OBJECT-GROUP
  OBJECTS {
    energyReactivePower,
    energyApparentPower,
    energyPhaseAngle,
    energyPhaseAnglePrecision
  }
  STATUS      current
  DESCRIPTION
    "The group of object for reporting details of
    AC power measurement."
  ::= { energyGroups 3 }

energyReactivePowerGroup OBJECT-GROUP
  OBJECTS {
    energyReactivePower
  }
  STATUS      current
  DESCRIPTION
    "The group of object for reporting the reactive power
    measured for AC supply."
  ::= { energyGroups 4 }

energyApparentPowerGroup OBJECT-GROUP
  OBJECTS {
    energyApparentPower
  }
  STATUS      current
  DESCRIPTION
    "The group of object for reporting the apparent power
    measured for AC supply."
  ::= { energyGroups 5 }

energyPhaseAngleGroup OBJECT-GROUP
  OBJECTS {
    energyPhaseAngle,
    energyPhaseAnglePrecision
  }
  STATUS      current
  DESCRIPTION
    "The group of object for reporting the phase angler
    measured for AC supply."
  ::= { energyGroups 6 }

END
```

8.3. Battery MIB

```
BATTERY-MIB DEFINITIONS ::= BEGIN
```

```
IMPORTS
```

```
    MODULE-IDENTITY, OBJECT-TYPE, NOTIFICATION-TYPE,  
    mib-2, Integer32, Unsigned32, Counter32  
        FROM SNMPv2-SMI -- RFC2578  
    DateAndTime  
        FROM SNMPv2-TC -- RFC2579  
    MODULE-COMPLIANCE, OBJECT-GROUP, NOTIFICATION-GROUP  
        FROM SNMPv2-CONF -- RFC2580  
    entPhysicalIndex  
        FROM ENTITY-MIB; -- RFC4133
```

```
batteryMIB MODULE-IDENTITY
```

```
    LAST-UPDATED "201001291200Z" -- 29 January 2010  
    ORGANIZATION "IETF OPSAWG Working Group"  
    CONTACT-INFO  
        "General Discussion: opsawg@ietf.org  
        To Subscribe: https://www.ietf.org/mailman/listinfo/opsawg  
        Archive: http://www.ietf.org/mail-archive/web/opsawg
```

```
    Co-editor:
```

```
        Juergen Quittek  
        NEC Europe Ltd.  
        NEC Laboratories Europe  
        Kurfuersten-Anlage 36  
        69115 Heidelberg  
        Germany  
        Tel: +49 6221 4342-115  
        Email: quittek@neclab.eu
```

```
    Co-editor:
```

```
        Thomas Dietz  
        NEC Europe Ltd.  
        NEC Laboratories Europe  
        Kurfuersten-Anlage 36  
        69115 Heidelberg  
        Germany  
        Phone: +49 6221 4342-128  
        Email: Thomas.Dietz@neclab.eu"
```

```
DESCRIPTION
```

```
    "This MIB module defines a set of objects for monitoring  
    batteries of networked devices and of their components.
```

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

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

This version of this MIB module is part of RFC yyyy; see the RFC itself for full legal notices."

-- replace yyyy with actual RFC number & remove this notice

-- Revision history

REVISION "201001291200Z" -- 29 January 2010
DESCRIPTION

"Initial version, published as RFC yyyy."

-- replace yyyy with actual RFC number & remove this notice

::= { mib-2 zzz }

-- zzz to be assigned by IANA.

-- *****
-- Top Level Structure of the MIB module
-- *****

batteryNotifications OBJECT IDENTIFIER ::= { batteryMIB 0 }
batteryObjects OBJECT IDENTIFIER ::= { batteryMIB 1 }
batteryConformance OBJECT IDENTIFIER ::= { batteryMIB 2 }

=====
-- 1. Object Definitions
=====

-- 1.1. Battery Table

batteryTable OBJECT-TYPE
SYNTAX SEQUENCE OF BatteryEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
"This table provides information on batteries in networked devices. It is designed as a sparse augment of the entPhysicalTable defined in the ENTITY-MIB module and assumes that each battery is represented by an individual row in the entPhysicalTable with an individual value for the index entPhysicalIndex.

Entries appear in this table only for entities that represent a battery. An entry in this table SHOULD be created at the same time as the associated entPhysicalEntry. An entry SHOULD be destroyed if the associated entPhysicalEntry is destroyed."

```
::= { batteryObjects 1 }
```

```
batteryEntry OBJECT-TYPE
```

```
SYNTAX      BatteryEntry
```

```
MAX-ACCESS  not-accessible
```

```
STATUS      current
```

```
DESCRIPTION
```

```
"An entry providing information on a battery."
```

```
INDEX { entPhysicalIndex } -- SPARSE-AUGMENTS
```

```
::= { batteryTable 1 }
```

```
BatteryEntry ::=
```

```
SEQUENCE {
```

batteryType	INTEGER,
batteryTechnology	INTEGER,
batteryNominalVoltage	Unsigned32,
batteryNumberOfCells	Unsigned32,
batteryNominalCapacity	Unsigned32,
batteryRemainingCapacity	Unsigned32,
batteryChargingCycleCount	Counter32,
batteryLastChargingCycleTime	DateAndTime,
batteryState	INTEGER,
batteryCurrentCharge	Unsigned32,
batteryCurrentChargePercentage	Unsigned32,
batteryCurrentVoltage	Unsigned32,
batteryCurrentCurrent	Integer32,
batteryLowAlarmPercentage	Unsigned32,
batteryLowAlarmVoltage	Unsigned32,
batteryReplacementAlarmCapacity	Unsigned32,
batteryReplacementAlarmCycles	Unsigned32

```
}
```

```
batteryType OBJECT-TYPE
```

```
SYNTAX      INTEGER {
                    primary(1),
                    rechargeable(2),
                    capacitor(3),
                    other(4),
                    unknown(5)
                }
```

```
MAX-ACCESS  read-only
```

```
STATUS      current
```

```
DESCRIPTION
```

"This object indicates the type of battery. It distinguishes between one-way primary batteries, rechargeable secondary batteries and capacitors which are not really batteries but often used in the same way as a battery.

The value other(4) can be used if the battery type is known but none of the ones above. Value unknown(5) is to be used if the type of battery cannot be determined."

::= { batteryEntry 1 }

batteryTechnology OBJECT-TYPE

```
SYNTAX      INTEGER {
                zincCarbon(1),
                zincChloride(2),
                oxyNickelHydroxide(3),
                lithiumCopper(4),
                lithiumIron(5),
                lithiumManganese(6),
                zincAir(7),
                silverOxide(8),
                alkaline(9),
                leadAcid(10),
                nickelCadmium(12),
                nickelMetalHybride(13),
                nickelZinc(14),
                lithiumIon(15),
                lithiumPolymer(16),
                doubleLayerCapacitor(17),
                other(18),
                unknown(19)
            }
```

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the technology used by the battery. Values 1-8 are primary battery technologies, values 10-16 are rechargeable battery technologies and value alkaline(9) is used for primary batteries as well as for rechargeable batteries.

The value other(18) can be used if the battery type is known but none of the ones above. Value unknown(19) is to be used if the type of battery cannot be determined."

::= { batteryEntry 2 }

batteryNominalVoltage OBJECT-TYPE

SYNTAX Unsigned32

UNITS "millivolt"

MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object provides the nominal voltage of the battery
in units of millivolt (mV).

Note that the nominal voltage is a constant value and
typically different from the actual voltage of the battery.

A value of 0 indicates that the nominal voltage is unknown."
 ::= { batteryEntry 3 }

batteryNumberOfCells OBJECT-TYPE

SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object indicates the number of cells contained in the
battery.

A value of 0 indicates that the number of cells is unknown."
 ::= { batteryEntry 4 }

batteryNominalCapacity OBJECT-TYPE

SYNTAX Unsigned32
UNITS "milliampere hours"
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object provides the nominal capacity of the battery
in units of milliampere hours (mAh).

Note that the nominal capacity is a constant value and
typically different from the actual capacity of the battery.

A value of 0 indicates that the nominal capacity is unknown."
 ::= { batteryEntry 5 }

batteryRemainingCapacity OBJECT-TYPE

SYNTAX Unsigned32
UNITS "milliampere hours"
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object provides the ACTUAL REMAINING capacity of the
battery in units of milliampere hours (mAh).

Note that the actual capacity needs to be measured and is

typically an estimate based on observed discharging and charging cycles of the battery.

A value of 'ffffffff'H indicates that the actual capacity cannot be determined."

::= { batteryEntry 6 }

batteryChargingCycleCount OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object indicates the number of charging cycles that that the battery underwent. Please note that the precise definition of a recharge cycle varies for different kinds of batteries and of devices containing batteries.

For batteries of type primary(1) the value of this object is always 0.

A value of 'ffffffff'H indicates that the number of charging cycles cannot be determined."

::= { batteryEntry 7 }

batteryLastChargingCycleTime OBJECT-TYPE

SYNTAX DateAndTime

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The date and time of the last charging cycle. The value '0000000000000000'H is returned if the battery has not been charged yet or if the last charging time cannot be determined.

For batteries of type primary(1) the value of this object is always '0000000000000000'H."

::= { batteryEntry 8 }

batteryState OBJECT-TYPE

SYNTAX INTEGER {
 full(1),
 partiallyCharged(2),
 empty(3),
 charging(4),
 discharging(5),
 unknown(6)
}

MAX-ACCESS read-only

STATUS current
DESCRIPTION
"This object indicates the current state of the battery.
Value full(1) indicates a full battery with a capacity
given by object batteryRemainingCapacity. Value empty(3)
indicates a battery that cannot be used for providing
electric power before charging it. Value partiallyCharged(2)
is provided if the battery is neither empty nor full and if
no charging or discharging is in progress. Charging or
discharging of the battery is indicated by values charging(3)
or discharging(4), respectively.

Value unknown(6) is to be used if the state of the battery
cannot be determined."
 ::= { batteryEntry 9 }

batteryCurrentCharge OBJECT-TYPE
SYNTAX Unsigned32
UNITS "milliampere hours"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
"This object provides the current charge of the battery
in units of milliampere hours (mAh).

Note that the current charge needs to be measured and is
typically an estimate based on observed discharging and
charging cycles of the battery.

A value of 'ffffffff'H indicates that the current charge
cannot be determined."
 ::= { batteryEntry 10 }

batteryCurrentChargePercentage OBJECT-TYPE
SYNTAX Unsigned32 (0..10000)
MAX-ACCESS read-only
STATUS current
DESCRIPTION
"This object provides the current charge of the battery
relative to the nominal capacity in units of a hundreds
of a percent.

-- Open issue:
-- Should it be the percentage of the nominal capacity
-- or of the current capacity?

Note that this value needs to be measured and is typically an estimate based on observed discharging and charging cycles of the battery.

A value of 'ffffffff'H indicates that the relative current charge cannot be determined."

::= { batteryEntry 11 }

batteryCurrentVoltage OBJECT-TYPE

SYNTAX Unsigned32

UNITS "millivolt"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object provides the current voltage of the battery in units of millivolt (mV).

A value of 'ffffffff'H indicates that the current voltage cannot be determined."

::= { batteryEntry 12 }

batteryCurrentCurrent OBJECT-TYPE

SYNTAX Integer32

UNITS "milliampere"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object provides the current charging or discharging current of the battery in units of milliampere (mA). Charging current is indicated by positive values, discharging current is indicated by negative values.

A value of '7ffffffff'H indicates that the current current cannot be determined."

::= { batteryEntry 13 }

batteryLowAlarmPercentage OBJECT-TYPE

SYNTAX Unsigned32 (0..10000)

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object provides the lower threshold value for object batteryCurrentChargePercentage. If the value of object batteryCurrentChargePercentage falls below this threshold, a low battery alarm will be raised. The alarm procedure may include generating a batteryLowNotification.

A value of 0 indicates that the no alarm will be raised for

any value of object batteryCurrentChargePercentage."
 ::= { batteryEntry 14 }

batteryLowAlarmVoltage OBJECT-TYPE

SYNTAX Unsigned32
UNITS "millivolt"
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object provides the lower threshold value for object batteryCurrentVoltage. If the value of object batteryCurrentVoltage falls below this threshold, a low battery alarm will be raised. The alarm procedure may include generating a batteryLowNotification.

A value of 0 indicates that the no alarm will be raised for any value of object batteryCurrentVoltage."

::= { batteryEntry 15 }

batteryReplacementAlarmCapacity OBJECT-TYPE

SYNTAX Unsigned32
UNITS "milliampere hours"
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object provides the lower threshold value for object batteryRemainingCapacity. If the value of object batteryRemainingCapacity falls below this threshold, a battery aging alarm will be raised. The alarm procedure may include generating a batteryAgingNotification.

A value of 0 indicates that the no alarm will be raised for any value of object batteryRemainingCapacity."

::= { batteryEntry 16 }

batteryReplacementAlarmCycles OBJECT-TYPE

SYNTAX Unsigned32
UNITS "milliampere hours"
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"This object provides the upper threshold value for object batteryChargingCycleCount. If the value of object batteryChargingCycleCount rises above this threshold, a battery aging alarm will be raised. The alarm procedure may include generating a batteryAgingNotification.

A value of 0 indicates that the no alarm will be raised for

```
any value of object batteryChargingCycleCount."
 ::= { batteryEntry 17 }
```

```
-----
-- 2. Notifications
-----
```

```
batteryLowNotification NOTIFICATION-TYPE
```

```
  OBJECTS      {
    batteryCurrentChargePercentage,
    batteryCurrentVoltage
  }
```

```
  STATUS      current
```

```
  DESCRIPTION
```

```
    "This notification can be generated when the current charge
    (batteryCurrentChargePercentage) or the current voltage
    (batteryCurrentVoltage) of the battery falls below a
    threshold defined by object batteryLowAlarmPercentage or
    object batteryLowAlarmVoltage, respectively."
```

```
 ::= { batteryNotifications 1 }
```

```
batteryAgingNotification NOTIFICATION-TYPE
```

```
  OBJECTS      {
    batteryRemainingCapacity,
    batteryChargingCycleCount
  }
```

```
  STATUS      current
```

```
  DESCRIPTION
```

```
    "This notification can be generated when the remaining
    capacity (batteryRemainingCapacity) falls below a threshold
    defined by object batteryReplacementAlarmCapacity
    or when the charging cycle count of the battery
    (batteryChargingCycleCount) exceeds the threshold defined
    by object batteryLowAlarmPercentage."
```

```
 ::= { batteryNotifications 2 }
```

```
-----
-- 3. Conformance Information
-----
```

```
batteryCompliances OBJECT IDENTIFIER ::= { batteryConformance 1 }
batteryGroups      OBJECT IDENTIFIER ::= { batteryConformance 2 }
```

```
-----
-- 3.1. Compliance Statements
-----
```

```
batteryCompliance MODULE-COMPLIANCE
  STATUS      current
  DESCRIPTION
    "The compliance statement for implementations of the
    POWER-STATE-MIB module.

    A compliant implementation MUST implement the objects
    defined in the mandatory group psmRequiredGroup."
  MODULE -- this module
  MANDATORY-GROUPS {
    batteryDescriptionGroup,
    batteryStatusGroup,
    batteryAlarmThresholdsGroup
  }
  GROUP      batteryNotificationsGroup
  DESCRIPTION
    "A compliant implementation does not have to implement
    the psmNotificationsGroup."
  ::= { batteryCompliances 1 }
```

```
-----
-- 3.2. MIB Grouping
-----
```

```
batteryDescriptionGroup OBJECT-GROUP
  OBJECTS {
    batteryType,
    batteryTechnology,
    batteryNominalVoltage,
    batteryNumberOfCells,
    batteryNominalCapacity
  }
  STATUS      current
  DESCRIPTION
    "A compliant implementation MUST implement the objects
    contained in this group."
  ::= { batteryGroups 1 }
```

```
batteryStatusGroup OBJECT-GROUP
  OBJECTS {
    batteryRemainingCapacity,
    batteryChargingCycleCount,
    batteryLastChargingCycleTime,
    batteryState,
    batteryCurrentCharge,
    batteryCurrentChargePercentage,
    batteryCurrentVoltage,
    batteryCurrentCurrent
```

```
    }
    STATUS          current
    DESCRIPTION
        "A compliant implementation MUST implement the objects
        contained in this group."
    ::= { batteryGroups 2 }

batteryAlarmThresholdsGroup OBJECT-GROUP
    OBJECTS {
        batteryLowAlarmPercentage,
        batteryLowAlarmVoltage,
        batteryReplacementAlarmCapacity,
        batteryReplacementAlarmCycles
    }
    STATUS          current
    DESCRIPTION
        "A compliant implementation MUST implement the objects
        contained in this group."
    ::= { batteryGroups 3 }

batteryNotificationsGroup NOTIFICATION-GROUP
    NOTIFICATIONS {
        batteryLowNotification,
        batteryAgingNotification
    }
    STATUS          current
    DESCRIPTION
        "A compliant implementation does not have to implement the
        notification contained in this group."
    ::= { batteryGroups 4 }
END
```

9. Security Considerations

There are no management objects defined in this MIB module that have a MAX-ACCESS clause of read-write and/or read-create. So, if this MIB module is implemented correctly, then there is no risk that an intruder can alter or create any management objects of this MIB module via direct SNMP SET operations.

Some of the readable objects in this MIB module (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 even GET and/or NOTIFY access to these objects and possibly to even encrypt the values of these objects when sending them over the network via SNMP. These are the tables and objects and their sensitivity/vulnerability:

- o This list is still to be done.

SNMP versions prior to SNMPv3 did not include adequate security. Even if the network itself is secure (for example by using IPsec), even then, there is no control as to who on the secure network is allowed to access and GET/SET (read/change/create/delete) the objects in this MIB module.

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 this MIB module is properly configured to give access to the objects only to those principals (users) that have legitimate rights to indeed GET or SET (change/create/delete) them.

10. IANA Considerations

The MIB modules in this document uses the following IANA-assigned OBJECT IDENTIFIER values recorded in the SMI Numbers registry:

Descriptor	OBJECT IDENTIFIER value
-----	-----
powerStateMIB	{ mib-2 xxx }
energyMIB	{ mib-2 yyy }
batteryMIB	{ mib-2 zzz }

Other than that this document does not impose any IANA considerations.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.

- [RFC2579] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Textual Conventions for SMIV2", STD 58, RFC 2579, April 1999.
- [RFC2580] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Conformance Statements for SMIV2", STD 58, RFC 2580, April 1999.
- [RFC4268] Chisholm, S. and D. Perkins, "Entity State MIB", RFC 4268, November 2005.
- [RFC3621] Berger, A. and D. Romascanu, "Power Ethernet MIB", RFC 3621, December 2003.
- [RFC3433] Bierman, A., Romascanu, D., and K. Norseth, "Entity Sensor Management Information Base", RFC 3433, December 2002.
- [RFC4133] Bierman, A. and K. McCloghrie, "Entity MIB (Version 3)", RFC 4133, August 2005.
- [I-D.quittek-power-monitoring-requirements] 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.
- [I-D.claise-energy-monitoring-mib] Claise, B., Chandramouli, M., Parello, J., and B. Schoening, "Power and Energy Monitoring MIB", draft-claise-energy-monitoring-mib-06 (work in progress), October 2010.
- [I-D.parello-eman-energy-aware-mib] Claise, B. and J. Parello, "Energy-aware Networks and Devices MIB", draft-parello-eman-energy-aware-mib-00 (work in progress), October 2010.

11.2. Informative References

- [RFC1628] Case, J., "UPS Management Information Base", RFC 1628, May 1994.
- [RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.

Authors' Addresses

Juergen Quittek (editor)
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-115
Email: quittek@nw.neclab.eu

Rolf Winter
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-121
Email: Rolf.Winter@nw.neclab.eu

Thomas Dietz
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-128
Email: Thomas.Dietz@nw.neclab.eu

Dominique Dudkowski
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-233
Email: Dominique.Dudkowski@nw.neclab.eu

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 25, 2011

J. Quittek, Ed.
R. Winter
T. Dietz
NEC Europe Ltd.
B. Claise
M. Chandramouli
Cisco Systems, Inc.
October 22, 2010

Requirements for Power Monitoring
draft-quittek-power-monitoring-requirements-02

Abstract

This memo discusses requirements for energy management, particularly for monitoring energy consumption and controlling power states of managed devices. This memo further shows that existing IETF standards are not sufficient for energy management and that energy management requires architectural considerations that are different from common other management functions.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 25, 2011.

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.

Table of Contents

1. Introduction	4
1.1. Energy management functions	4
1.2. Specific aspects of energy management	6
2. Scenarios and target devices	6
2.1. Scenario 1: Routers, switches, middleboxes, and hosts . .	6
2.2. Scenario 2: PoE sourcing equipment and PoE powered devices	7
2.3. Scenario 3: Power probes and Smart meters	7
2.4. Scenario 4: Mid-level managers	7
2.5. Scenario 5: Gateways to building networks	7
2.6. Scenario 6: Home energy gateways	8
2.7. Scenario 7: Data center devices	8
2.8. Scenario 8: Battery powered devices	8
3. Monitoring Requirements	8
3.1. Granularity of monitoring and control	8
3.2. Remote and Aggregated Monitoring	9
3.3. Accuracy	10
3.4. Required Information	10
3.4.1. Power State Monitoring	10
3.4.2. Energy Consumption Monitoring	11
3.4.3. Power Quality	11
3.4.4. Battery State Monitoring	12
4. Monitoring Models	12
5. Control Requirements	13
6. Existing Standards	13
6.1. Existing IETF Standards	13
6.1.1. ENTITY STATE MIB	13
6.1.2. ENTITY SENSOR MIB	14
6.1.3. UPS MIB	14
6.1.4. POWER ETHERNET MIB	15
6.1.5. LLDP MED MIB	15
6.2. Existing standards of other bodies	15
6.2.1. DMTF	15

7. Suggested Actions	16
8. Acknowledgements	16
9. Security Considerations	16
10. IANA Considerations	17
11. Informative References	17
Authors' Addresses	18

1. Introduction

With rising energy cost and with an increasing awareness of the ecological impact of running IT and networking equipment, energy management is becoming an additional basic requirement for network management frameworks and systems.

Different to other typical network management functions, energy management often extends its scope beyond devices with IP network interfaces. Requirements in this document do not fully cover all these networks, but they cover means for opening IP network management towards them.

In general, IETF Standards for energy management should be defined in such a way that they can be applied to several areas including

- o Communication networks and IT systems
- o Building networks
- o Home networks
- o Smart (power) grids

1.1. Energy management functions

The basic objective of energy management is operating communication networks and other equipment with a minimal amount of energy. An energy management system should provide means for reducing the power consumption of individual components of a network as well as of the whole network.

One approach to achieve this goal is setting all components to an operational state that results in lower energy consumption but still meets service level performance objectives. The sufficient performance level may vary over time and can depend on several factors. In principle, there are four basic types of power states for a component or for a whole system:

- o full power state
- o reduced power states (lower clock rate for processor, lower data rate on a link, etc.)
- o stand-by/sleep state (not functional, but immediately available)
- o power-off state (requiring significant time for becoming operational)

In actual implementations the number of power states and their properties vary a lot. Very simple devices may just have a full power and a power off state, while other devices may have a high number of different reduced power and sleep states.

While the general objective of energy management is quite clear, the way to attain that goal is often difficult. In many cases there is no way of reducing power consumption without the consequence of a

potential performance degradation. Then a trade-off needs to be dealt with between service level objectives and energy efficiency. In other cases a reduction of energy consumption can easily be achieved while still maintaining sufficient service level performance, for example, by switching components to lower power states when higher performance is not needed.

Network management systems can control such situations by implementing policies to achieve a certain degree of energy efficiency. In order to make policy decisions properly, information about the energy consumption of network components and sub-components in different power states is required. Often this information is acquired best through monitoring.

Monitoring operational power states and energy consumption is also useful for other energy management purposes including but not limited to

- o investigating power saving potential
- o evaluating the effectiveness of energy saving policies and measures
- o deriving, implementing, and testing power management strategies
- o accounting the total power consumption of a network element, a network, a service, or subcomponents of those

From the considerations described above the following basic management functions appear to be required for energy management:

- o monitoring power states of network elements and their subcomponents
- o monitoring actual power (energy consumption rate) of network elements and their subcomponents
- o monitoring (accumulated) energy consumption of network elements and their subcomponents
- o setting power states of network elements and their subcomponents
- o setting and enforcing power saving policies

Editorial note: With the extension to power state control and policy enforcement, the title of the draft does not anymore match the scope well. The name of the draft will be updated in a future revision.

It should be noted that monitoring energy consumption and power states itself is obviously not in itself a means to reduce the energy consumption of a device. In fact, it is likely to increase the power consumption of a device slightly. However, the acquired energy consumption and power state information is essential for defining energy saving policies and can be used as input to power state control loops that in total can lead to energy savings.

It should further be noted that active power control is complementary

(but essential) to other energy savings measures such as low power electronics, energy saving protocols (e.g. IEEE 802.3az), and energy-efficient device design (for example, stand-by and low-power modes for individual components of a device), and energy-efficient network architectures. Measurement of energy consumption may also provide input for developing these technologies.

1.2. Specific aspects of energy management

There are two aspects of energy management that make it different from other common network management functions. The first difference is that energy consumption is often measured remotely to the device under consideration. A reason for this is that today very few devices are instrumented with the hardware and software for measuring their own current power and accumulated energy consumption. Often power and energy for such devices is measured by other devices.

A common example is a Power over Ethernet (PoE) sourcing device that provides means for measuring provided power per port. If the device connected to a port is known, power and energy measurements for that device can be conducted by the PoE sourcing device. Another example is a smart power strip. Again, if it is known which devices are plugged into which outlets of the smart power strip, then the power strip can provide measured values for these devices.

The second difference is that often it is desirable to apply energy management also to networks and devices that do not communicate via IP, for example, in building networks where besides IP several other communication protocols are used. In these networks, it may be desirable that devices with IP interfaces report energy and power values for other devices. Reports may be based on measurements at the reporting device, similar to the PoE sourcing device and the smart strip. But reports may also be just relayed from non-IP communication to IP communication.

2. Scenarios and target devices

This section describes a selection of scenarios for the application of energy management. For each scenario a list of target devices is given in the headline, for which IETF energy management standards are needed.

2.1. Scenario 1: Routers, switches, middleboxes, and hosts

Power management of network devices is considered as a fundamental (basic first step) requirement. The devices listed in this scenario are some of the components of a communication network. For these

network devices, the chassis draws power from an outlet and feeds all its internal sub-components.

2.2. Scenario 2: PoE sourcing equipment and PoE powered devices

This scenario covers devices using Power over Ethernet (PoE). A PoE Power Sourcing Equipment (PSE), for example, a PoE switch, provides power to a PoE Powered Device (PD), for example, a PoE desktop phone. Here, the PSE provides means for controlling power supply (switching it on and off) and for monitoring actual power provided at a port to a specific PD.

2.3. Scenario 3: Power probes and Smart meters

Today, very few devices are equipped with sufficient instrumentation to measure their own actual power and accumulated energy consumption. Often external probes are connected to the power supply for measuring these properties for a single device or for a set of devices.

Homes, buildings, and data centers have smart meters that monitor and report accumulated power consumption of an entire home, a set of offices or a set of devices in data centers.

Power Distribution Unit (PDUs) attached to racks in data center and other smart power strips are evolving with smart meters and remote controllable power switches embedded for each socket.

2.4. Scenario 4: Mid-level managers

Sometimes it is useful to have mid-level managers that provide energy management functions not just for themselves but also for a set of associated devices. For example, a switch can provide energy management functions for all devices connected to its ports, even if these devices are not PoE PDs, but have their own power supply as, for example, PCs connected to the switch.

2.5. Scenario 5: Gateways to building networks

Due to the potential energy savings, energy management of buildings has received significant attention. There are gateway network elements to manage the multiple components of a building energy management network Heating Ventilating Air Conditioning (HVAC), lighting, electrical, fire and emergency systems, elevators etc. The gateway device provides power monitoring and control function for other devices in the building network.

2.6. Scenario 6: Home energy gateways

Home energy gateway can be used for energy management of a home. This gateway can manage the appliances (refrigerator, heating/cooling, washing machine etc.) and interface with the electrical grid. The gateway can implement policies based on demand and energy pricing from the grid.

2.7. Scenario 7: Data center devices

Energy efficiency of data centers has become a fundamental challenges of data center operation. Energy management is conducted on different aggregation levels, such as network level, Power Distribution Unit (PDU) level, and server level.

2.8. Scenario 8: Battery powered devices

Some devices have a battery as a back-up source of power. Given the finite capacity and lifetime of a battery, means for reporting the actual charge, age, and state of a battery are required.

3. Monitoring Requirements

3.1. Granularity of monitoring and control

Often it is desirable to switch off individual components of a device but not the entire device. The switch may need to continue serving a few ports (for example, the ports serving an email server or needed for server backup), but most other ports could be entirely switched off under some policies (for example at night or the weekend in an office).

As illustrated by this example, it is often desirable to monitor power state and energy consumption on a granularity level that is finer than just the device level. Monitoring should be available for individual components of devices, such as line cards, processor cards, hard drives, etc. For example, for IP routers the following list of views of a router gives an idea of components that potentially could be monitored and controlled individually:

- o Physical view: chassis (or stack), central control engine, line cards, service cards, etc.
- o Component view: CPU, ASICs, fans, power supply, ports (single ports and port groups), storage and memory
- o Feature view: L2 forwarding, L3 routing, security features, load balancing features, network management, etc.

- o Logical view: system, data-plane, control-plane, etc.
- o Relationship view: line cards, ports and the correlation between transmission speed and power consumption, relationship of system load and total power consumption

Instrumentation for measuring energy consumption of a device is typically more expensive than instrumentation for retrieving the devices power state. It may be a reasonable compromise in many cases to provide power state information for all individually switchable components of a device separately, while the energy consumption is only measured for the entire device.

3.2. Remote and Aggregated Monitoring

There are several ways power and energy consumption can be measured and reported. Measurements can be performed locally at the device that consumes energy or remotely by a device that has access to the power supply of another device.

Instrumentation for power and energy measurements at a device requires additional hardware. A cost-efficient alternative is measuring power and energy consumption aggregated for a set of devices, for example a PoE PSE reporting these values per port group instead of per port, or a power distribution unit that reports the values for all connected devices instead of per socket.

If aggregated measurement is conducted, it is obvious that reporting provides aggregated values. but aggregated reporting can also be combined with local measurements. A managed node may act as mid-level manager or protocol converter for several devices that measure power consumption by themselves, for example a home gateway or a gateway to building networks. In this case, the reporting node may choose to report for each device individual values or aggregated values from groups of devices that transmitted their power and energy consumption values to the reporting node.

Often it is sufficient and more cost efficient having a single device measuring and providing power state and energy consumption information not just for itself but also for several further devices that are in some way attached to it. If the measuring and reporting device has access to individual power supply lines for each device, then it can measure energy consumption per device. If it only has access to a joint power supply for several devices, then it will measure aggregated values.

One example for the first case is a switch acting as power sourcing equipment for several IP phones using Power over Ethernet (PoE). The switch can measure the power consumption for each phone individually

at the port the phone is connected to or it measures aggregated values per port group for a set of devices.. The phones do not need to provide means for energy consumption measurement and reporting by themselves.

Another example is a smart meter that just measures and reports the energy transmitted through attached electric cables. Such a smart meter can be used to monitor energy consumption of an individual device if connected to the devices' individual power supply. But in many common cases it measures the aggregated energy consumption of several devices, for example, as part of an uninterruptible power supply (UPS) that serves several devices at a single power cord, or as a smart electric meter for a set of machines in a rack, in an office building or at a residence.

3.3. Accuracy

Depending on how power and energy consumption values are obtained the confidence in the reported value and its accuracy may vary. Managed nodes reporting values concerning themselves or other devices should qualify the confidence in reported values and quantify the accuracy of measurements. For accuracy reporting, the accuracy classes specified in IEC 61850 should be considered.

3.4. Required Information

This section lists requirements for information to be retrieved. Because of the different nature of power state monitoring and energy consumption monitoring, these are discussed separately. In addition, a section on battery monitoring is included which again comes with a set of very different requirements.

Not all of the individual requirements listed in subsections below are equally relevant. A classification into 'required' and 'optional' is still in progress.

3.4.1. Power State Monitoring

The power state of a device or component typically can only have a small number of discrete values such as, for example, full power, low power, standby, hibernating, off. However, some of these states may have one or more sub-states or state parameters. For example, in low power state, a reduced clock rate may be set to a large number of different values. For the device power state, the following information is considered to be relevant:

- o the current state - the time of the last change
- o the cause for the last transition
- o time to transit from one stage to another
- o the total time spent in each state
- o the duration of the last period spent in each state
- o the number of transitions to each state
- o the current power source

For some network management tasks it may be desirable to receive notifications from devices when components or the entire device change their power state.

3.4.2. Energy Consumption Monitoring

Independent of the power state, energy consumption of a device or a device's component is a quantity for which the value may change continuously. Therefore, the information that needs to be retrieved concerning this quantity is quite different:

- o the current real power (energy consumption rate) averaged over a short time interval
- o total energy consumption
- o energy consumption since the last report or for the last configured time interval
- o total energy consumption per power state
- o energy consumption per power state since the last report

For some network management tasks it may be desirable to receive notifications from devices when the current power consumption of a component or of the entire device exceeds or falls below certain thresholds.

Energy consumption of a device or a device's component is a quantity for which the value may change continuously. For some network management tasks it is required to measure the power over time with a relatively high time resolution. In such a case not just single values for the current power of a component is needed, but a series of power values reporting on consecutive time intervals.

In order to put measured data into perspective, the precision of the measured data, i.e. the potential error in the measured data, needs to be known as well.

3.4.3. Power Quality

In addition to the quantity of power or energy, also power quality should be reported according to IEC 62053-22 and IEC 60044-1.

3.4.4. Battery State Monitoring

An increasing number of networked devices are expected to be battery powered. This includes e.g. smart meters that report meter readings and are installed in places where external power supply is not always possible or costly. But also other devices might have internal/external batteries to power devices for short periods of time when the main power fails, to power parts of the device when the main device is switches off etc. Knowing the state of these batteries is important for the operation of these devices and includes information such as:

- o the current charge of the battery
- o the age of the battery
- o the state of the battery (e.g. being re-charged)
- o last usage of the battery
- o maximum energy provided by the battery

It is possible for devices that are only battery-powered to send notifications when the current battery charge has dropped below a certain threshold in order to inform the management system of needed replacement. The same applies for the age of a battery.

4. Monitoring Models

Monitoring of power states and energy consumption can be performed in pull mode (for example, SNMP GET [RFC3410]) or in push mode (for example SNMP notification [RFC3410], Syslog [RFC5424], or IPFIX [RFC5101]).

Pull mode monitoring is often easier to handle for a network management systems, because it can determine when it gets certain information from a certain device. However, the overhead of pull model monitoring is typically higher than for push model monitoring, particularly when large numbers of values are to be collected, such as time series of power values.

In such cases, push model monitoring may be preferable with a device sending a data stream of values without explicit request for each value from the network management system. For notifications on events, only the push model is considered to be appropriate.

Applying these considerations to the required information leads to the conclusion that most of the information can appropriately be reported using the pull model. The only exceptions are notifications on power state changes and high volume time series of energy consumption values.

5. Control Requirements

To realize the envisioned benefits of energy savings, just monitoring power states and energy consumption would not be sufficient. Energy efficiency can be realized only by setting the network entities or components to energy saving power states when appropriate.

With means for power state control, energy saving policies and control loops can be realized. Policies may, for example, define different power state settings based on the time-of-day. Control loops may, for example, change power states based on actual network load.

Trivially, all entities being subject of energy management should have at least two power states, such as "on" and "off" or "on" and "sleep" to be set. In many cases, it may be desirable to have more operational ("on" mode) and non-operational ("off"/"sleep" mode) power states. This applies particularly to devices with a lot of configuration parameters that influence their energy consumption. Examples for specifications of power states of managed devices can be found in the Advanced Configuration and Power Interface (ACPI) [ACPI.R30b] or the DMTF Power State Management Profile [DMTF.DSP1027].

6. Existing Standards

This section analyzes existing standards for energy consumption and power state monitoring. It shows that there are already several standards that cover some part of the requirements listed above, but even all together they do not cover all of the requirements for energy management.

6.1. Existing IETF Standards

There are already RFCs available that address a subset of the requirements.

6.1.1. ENTITY STATE MIB

RFC 4268 [RFC4268] defines the ENTITY STATE MIB module. Implementations of this module provide information on entities including the standby status (hotStandby, coldStandby, providingService), the operational status (disabled, enabled, testing), the alarm status (underRepair, critical, major, minor, warning), and the usage status (idle, active, busy). This information is already useful as input to policy decisions and for other network monitoring tasks. However, the number of states would

cover only a small subset of the requirements for power state monitoring and it does not provide means for energy consumption monitoring. For associating the provided information to specific components of a device, the ENTITY STATE MIB module makes use of the means provided by the ENTITY MIB module [RFC4133]. Particularly, it uses the entPhysicalIndex for identifying entities.

The standby status provided by the ENTITY STATE MIB module is related to power states required for energy management, but the number of states is too restricted for meeting all energy management requirements. For energy management several more power states are required, such as different sleep and operational states as defined by the Advanced Configuration and Power Interface (ACPI) [ACPI.R30b] or the DMTF Power State Management Profile [DMTF.DSP1027].

6.1.2. ENTITY SENSOR MIB

RFC 3433 [RFC3433] defines the ENTITY SENSOR MIB module. Implementations of this module offer a generic way to provide data collected by a sensor. A sensor could be an energy consumption meter delivering measured values in Watt. This could be used for reporting current power of a device and its components. Furthermore, the ENTITY SENSOR MIB can be used to retrieve the precision of the used power meter.

Similar to the ENTITY STATE MIB module, the ENTITY SENSOR MIB module makes use of the means provided by the ENTITY MIB module [RFC4133] for relating provided information to components of a device.

However, there is no unit available for reporting energy quantities, such as, for example, watt seconds or kilowatt hours, and the ENTITY SENSOR MIB module does not support reporting accuracy of measurements according to the IEC / ANSI accuracy classes, which are commonly in use for electric power and energy measurements. The ENTITY SENSOR MIB modules only provides a coarse-grained method for indicating accuracy by stating the number of correct digits of fixed point values.

6.1.3. UPS MIB

RFC 1628 [RFC1628] defines the UPS MIB module. Implementations of this module provide information on the current real power of devices attached to an uninterruptible power supply (UPS) device. This application would require identifying which device is attached to which port of the UPS device.

UPS MIB provides information on the state of the UPS network. The MIB module contains several variables identify the UPS entity (name,

model,...), the battery state, to characterize the input load to the UPS, to characterize the output from the UPS, to indicate the various alarm events. The measurements of power in UPS MIB are in Volts, Amperes and Watts. The units of power measurement are RMS volts, RMS Amperes and are not based on Entity-Sensor MIB [RFC3433].

6.1.4. POWER ETHERNET MIB

Similar to the UPS MIB, implementations of the POWER ETHERNET MIB module defined in RFC3621 [RFC3621] provide information on the current energy consumption of the devices that receive Power over Ethernet (PoE). This information can be retrieved at the power sourcing equipment. Analogous to the UPS MIB, it is required to identify which devices are attached to which port of the power sourcing equipment.

The POWER ETHERNET MIB does not report power and energy consumption on a per port basis, but can report aggregated values for groups of ports. It does not use objects of the ENTITY MIB module for identifying entities, although this module existed already when the POWER ETHERNET MIB modules was standardized.

6.1.5. LLDP MED MIB

The Link Layer Discovery Protocol (LLDP) defined in IEEE 802.1ab is a data link layer protocol used by network devices for advertising of their identities, capabilities, and interconnections on a LAN network. The Media Endpoint Discovery (MED) (ANSI/TIA-1057) is an enhancement of LLDP known as LLDP-MED. The LLDP-MED enhancements specifically address voice applications. LLDP-MED covers 6 basic areas: capabilities discovery, LAN speed and duplex discovery, network policy discovery, location identification discovery, inventory discovery, and power discovery.

6.2. Existing standards of other bodies

6.2.1. DMTF

The DMTF has defined a power state management profile [DMTF.DSP1027] that is targeted at computer systems. It is based on the DMTF's Common Information Model (CIM) and rather a device profile than an actual energy consumption monitoring standard.

The power state management profile is used to describe and to manage the power state of computer systems. This includes e.g. means to change the power state of a device (e.g. to shutdown the device) which is an aspect of but not sufficient for active energy management.

7. Suggested Actions

Based on the analysis of requirements in Section 3 and the discussion of monitoring models in Section 4 this memo proposes to develop a standard for pull-based monitoring of power states, energy consumption, and battery states. The analysis of existing MIB modules in the previous section shows that they are not sufficient to meet the requirements discussed in Section 3.

As a consequence, it suggested to develop one or more MIB modules for this purpose. Such MIB modules could also cover push-based reporting of power state changes using SNMP notifications. The only aspect that would not be covered well by a MIB/SNMP solution is the reporting of large time series of energy consumption values. For this purpose SNMP does not appear to be an optimal choice. Particularly for supporting smart meter functionality, a push-based protocol appears to be more appropriate. Within the IP protocol family the Syslog and IPFIX protocols seem to be the most suitable candidates. There are more standard protocols with the capability to transfer measurement series, for example, DIAMETER, but these protocols are designed and well suited for other application areas than network monitoring.

Comparing the two candidates (Syslog and IPFIX), IPFIX seems to be the better suited one. While Syslog is optimized for the transmission of text messages, IPFIX is better equipped for transmitting sequences of numerical values. Encoding numerical values into syslog is well feasible, see, for example, the mapping of SNMP notifications to Syslog messages in [RFC5675], but IPFIX provides better means. With the extensible IPFIX information model [RFC5102] no protocol extension would be required for transmitting energy consumption information. Only a set of new information elements would need to be registered at IANA. However, this memo suggests that the definition of such information elements should be conducted within the IETF and they should be documented in a standards track RFC.

8. Acknowledgements

The authors would like to thank Ralf Wolter, for his first essay on this draft.

9. Security Considerations

This memo currently does not impose any security considerations.

10. IANA Considerations

This memo has no actions for IANA..

11. Informative References

- [RFC4268] Chisholm, S. and D. Perkins, "Entity State MIB", RFC 4268, November 2005.
- [RFC3621] Berger, A. and D. Romascanu, "Power Ethernet MIB", RFC 3621, December 2003.
- [RFC1628] Case, J., "UPS Management Information Base", RFC 1628, May 1994.
- [RFC3433] Bierman, A., Romascanu, D., and K. Norseth, "Entity Sensor Management Information Base", RFC 3433, December 2002.
- [RFC4133] Bierman, A. and K. McCloghrie, "Entity MIB (Version 3)", RFC 4133, August 2005.
- [RFC5101] Claise, B., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", RFC 5101, January 2008.
- [RFC5102] Quittek, J., Bryant, S., Claise, B., Aitken, P., and J. Meyer, "Information Model for IP Flow Information Export", RFC 5102, January 2008.
- [RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.
- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.
- [RFC5675] Marinov, V. and J. Schoenwaelder, "Mapping Simple Network Management Protocol (SNMP) Notifications to SYSLOG Messages", RFC 5675, October 2009.
- [ACPI.R30b] Hewlett-Packard Corporation, Intel Corporation, Microsoft Corporation, Phoenix Corporation, and Toshiba Corporation, "Advanced Configuration and Power Interface Specification, Revision 3.0b", October 2006.
- [DMTF.DSP1027] Dasari (ed.), R., Davis (ed.), J., and J. Hilland (ed.),

"Power State Management Profile", September 2008.

Authors' Addresses

Juergen Quittek (editor)
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-115
Email: quittek@neclab.eu

Rolf Winter
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-121
Email: Rolf.Winter@neclab.eu

Thomas Dietz
NEC Europe Ltd.
NEC Laboratories Europe
Network Research Division
Kurfuersten-Anlage 36
Heidelberg 69115
DE

Phone: +49 6221 4342-128
Email: Thomas.Dietz@neclab.eu

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Degem 1831
BE

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

Mouli Chandramouli
Cisco Systems, Inc.
Sarjapur Outer Ring Road
Bangalore,
IN

Phone: +91 80 4426 3947
Email: moulchan@cisco.com

