

CoRE Working Group  
Internet-Draft  
Intended status: Informational  
Expires: January 7, 2016

T. Zotti  
Philips Research  
P. van der Stok  
Consultant  
E. Dijk  
Philips Research  
July 6, 2015

Sleepy CoAP Nodes  
draft-zotti-core-sleepy-nodes-03

## Abstract

6LoWPAN networks rely on application protocols like CoAP to enable RESTful communications in constrained environments. Many of these networks make use of "Sleepy Nodes": battery powered devices that switch off their (radio) interface during most of the time to conserve battery energy. As a result of this, Sleepy Nodes cannot be reached most of the time. This fact prevents using normal communication patterns as specified in the CoRE group, since the server-model is not applicable to these devices. This document discusses and specifies an architecture to support Sleepy Nodes such as battery-powered sensors in 6LoWPAN networks with the goal of guiding and stimulating the discussion on Sleepy Nodes support for CoAP in the CoRE WG.

## 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 January 7, 2016.

Internet-Draft

Sleepy Nodes

July 2015

## Copyright Notice

Copyright (c) 2015 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">1.1.</a>	Problem statement . . . . .	<a href="#">3</a>
<a href="#">1.2.</a>	Assumptions . . . . .	<a href="#">4</a>
<a href="#">1.3.</a>	Requirements Language . . . . .	<a href="#">4</a>
<a href="#">1.4.</a>	Acronyms . . . . .	<a href="#">5</a>
<a href="#">2.</a>	Use cases and architecture . . . . .	<a href="#">5</a>
<a href="#">2.1.</a>	Node interactions and use cases . . . . .	<a href="#">6</a>
<a href="#">2.2.</a>	Architecture . . . . .	<a href="#">9</a>
<a href="#">2.3.</a>	Example contents . . . . .	<a href="#">10</a>
<a href="#">3.</a>	Interactions involving Resource Directory . . . . .	<a href="#">10</a>
<a href="#">4.</a>	Synchronize interface . . . . .	<a href="#">11</a>
<a href="#">4.1.</a>	Sleepy node discovers proxy . . . . .	<a href="#">11</a>
<a href="#">4.2.</a>	Registration at a Proxy . . . . .	<a href="#">12</a>
<a href="#">4.3.</a>	De-registration at a Proxy . . . . .	<a href="#">14</a>
<a href="#">4.4.</a>	Initialization of delegated resource . . . . .	<a href="#">15</a>
<a href="#">4.5.</a>	Sleepy Node updates delegated resource at Proxy . . . . .	<a href="#">17</a>
<a href="#">4.6.</a>	Sleepy Node READs resource updates from Proxy . . . . .	<a href="#">17</a>
<a href="#">5.</a>	Delegate Interface . . . . .	<a href="#">18</a>
<a href="#">5.1.</a>	Discovering Endpoint discovers Sleepy Node at Proxy . . . . .	<a href="#">18</a>
<a href="#">5.2.</a>	Proxy REPORTs events to Endpoint . . . . .	<a href="#">19</a>
<a href="#">5.3.</a>	A Node WRITES to Sleepy Node via Proxy . . . . .	<a href="#">21</a>
<a href="#">5.4.</a>	A Node READs information from Sleepy Node via Proxy . . . . .	<a href="#">22</a>
<a href="#">6.</a>	Direct Interface . . . . .	<a href="#">22</a>
<a href="#">6.1.</a>	Sleepy Node REPORTs events directly to Destination Node . . . . .	<a href="#">22</a>
<a href="#">6.2.</a>	A Sleepy Node READs information from a Server Node . . . . .	<a href="#">23</a>
<a href="#">7.</a>	Realization with PubSub server . . . . .	<a href="#">23</a>

<a href="#">8.</a>	IANA Considerations . . . . .	<a href="#">23</a>
<a href="#">9.</a>	Security Considerations . . . . .	<a href="#">23</a>
<a href="#">10.</a>	Acknowledgements . . . . .	<a href="#">24</a>
<a href="#">11.</a>	Changelog . . . . .	<a href="#">24</a>
<a href="#">12.</a>	References . . . . .	<a href="#">24</a>

<a href="#">12.1.</a>	Normative References . . . . .	<a href="#">24</a>
<a href="#">12.2.</a>	Informative References . . . . .	<a href="#">25</a>
	Authors' Addresses . . . . .	<a href="#">25</a>

## [1.](#) Introduction

6LoWPAN networks rely on application protocols such as CoAP to enable RESTful communications in constrained environments. Many of these networks feature "Sleepy Nodes": battery-powered nodes which switch on/off their communication interface to conserve battery energy. As a result of this, Sleepy Nodes cannot be reached most of the time. This fact prevents using normal communication patterns as specified by the CoRE group, since the server model is clearly not applicable to the most energy constrained devices.

This document discusses and specifies an architecture to support Sleepy Nodes such as battery-powered sensors in 6LoWPAN networks. The proposed solution makes use of a Proxy Node to which a Sleepy Node delegates part of its communication tasks while it is not accessible in the 6LoWPAN network. Direct interactions between Sleepy Nodes and non-Sleepy Nodes are only possible, when the Sleepy Node initiates the communication.

Earlier related documents treating the Sleepy node subject are the CoRE mirror server [[I-D.vial-core-mirror-server](#)] and the Publish-Subscribe in the Constrained Application Protocol (CoAP) [[I-D.koster-core-coap-pubsub](#)]. Both documents describe the interfaces to the proxy accompanying the Sleepy node. Both make use of the observe option discussed in [[I-D.ietf-core-observe](#)]. This document describes the roles of the nodes communicating with the Sleepy node and/or its proxy. As such it contributes to understanding how well the other proposals support the operation of the Sleepy nodes in a building control context.

The issues that need to be addressed to provide support for Sleepy Nodes in 6LoWPAN networks are summarized in [Section 1.1](#). [Section 2](#)

provides a set of use case descriptions that introduce communication patterns to be used in home and building control scenarios.

[Section 3](#), [Section 4](#), [Section 5](#), and [Section 6](#) specify interfaces to support each of these scenarios. Many interface specifications and examples are taken over from [[I-D.vial-core-mirror-server](#)].

### [1.1](#). Problem statement

During typical operation, a Sleepy Node has its radio disabled and the CPU may be in a sleeping state. If an external event occurs (e.g. person walks into the room activating a presence sensor), the CPU and radio are powered back on and they send out an event message

to another node, or to a group of nodes. After sending this message, the radio and CPU are powered off again, and the Sleepy Node sleeps until the next external event or until a predefined time period has passed. The main problems when introducing Sleepy Nodes into a 6LoWPAN network are as follows:

Problem 1: How to contact a Sleepy Node that has its radio turned off most of the time for:

- Writing configuration settings.
- Reading out sensor data, settings or log data.
- Configuring additional event destination nodes or node groups.

Problem 2: How to discover a Sleepy Node and its services, while the node is asleep:

- Direct node discovery (CoAP GET `/.well-known/core` as defined in [[RFC7252](#)]) does not find the node with high probability.
- Mechanisms may be needed to provide, as the result of node discovery, the IP address of a Proxy instead of the IP address of the node directly.

Problem 3: How a Sleepy Node can convey data to a node or groups of nodes, with good reliability and minimal energy consumption.

### [1.2](#). Assumptions

The solution architecture specified here assumes that a Sleepy Node has enough energy to perform bidirectional communication during its normal operational state. This solution may be applicable also to extreme low-power devices such as solar powered sensors as long as they have enough energy to perform commissioning and the initial registration steps. These installation operations may require, in some cases, an additional source of power. Since a Sleepy Node is unreachable for relatively long periods of times, the data exchanges in the interaction model are always initiated by a Sleepy Node when its sleep period ends.

### [1.3.](#) Requirements Language

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](#)].

This document assumes readers are familiar with the terms and concepts discussed in [[RFC7252](#)], [[RFC5988](#)], [[I-D.ietf-core-resource-directory](#)], [[I-D.ietf-core-interfaces](#)], [[I-D.ietf-core-observe](#)] and [[I-D.vial-core-mirror-server](#)].

In addition, this document makes use of the following additional terminology:

**Sleepy Node:** a battery-powered node which does the on/off switching of its communication interface with the purpose of conserving battery energy

**Sleeping/Asleep:** A Sleepy Node being in a "sleeping state" i.e. its network interface is switched off and a Sleepy Node is not able to send or receive messages.

**Awake/Not Sleeping:** A Sleepy Node being in an "awake state" i.e. its network interface is switched on and the Sleepy Node is able to send or receive messages.

**Wake up reporting duration:** the duration between a wake up from a

Sleepy Node and the next wake up and report of the same Node.

Proxy: any node that is configured to, or selected to, perform communication tasks on behalf of one or more Sleepy Nodes.

Regular Node: any node in the network which is not a Proxy or a Sleepy Node.

#### [1.4.](#) Acronyms

This Internet-Draft contains the following acronyms:

DTLS: Datagram Transport Layer Security

EP: Endpoint

MC: Multicast

RD: Resource Directory

## [2.](#) Use cases and architecture

To describe the application viewpoint of the solution, we introduce some example scenarios for the various interactions shown in Figure 1. The figure assigns the following roles taken up by a regular node:

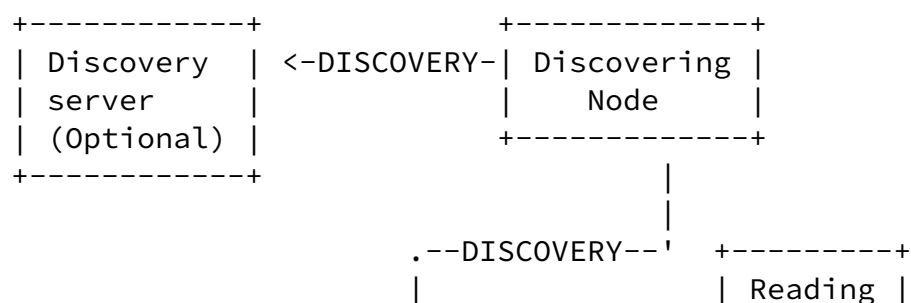
- o Reading Node: any regular node that reads information from the Sleepy Node.
- o Configuring Node: any regular node that writes information/configuration into Sleepy Node(s). Examples of configuration are new thresholds for a sensor or a new value for the wake-up cycle time.
- o Discovering Node: any regular node that performs discovery of the nodes in a network, including Sleepy Nodes.
- o Destination Node: any regular node or node in a group that receives a message that is generated by the Sleepy Node.
- o Server Node: an optional server that the Sleepy Node knows about,

or is told about, which is used to fetch  
information/configuration/firmware updates/etc.

- o Discovery Server: an optional server that enables nodes to discover all the devices in the network, including Sleepy Nodes, and query their capabilities. For example, a Resource Directory server as defined in [[I-D.ietf-core-resource-directory](#)] or a DNS-SD server as defined in [[RFC6763](#)]. For the rest of this document the discovery server is a Resource Directory. Specifically, the functionalities of the Resource Directory related to the architecture presented in this Internet-Draft are described in more details in [Section 3](#).

The term delegated resource is to indicate the copy at the Proxy of a resource present in the Sleepy Node.

### 2.1. Node interactions and use cases






Figure 1: Interaction model for Sleepy Nodes in 6LowPAN networks

The interactions visualized in the figure are discussed and motivated with their use cases:

**DISCOVERY Interaction:** a Discovering Node discovers Sleepy Node(s) via Proxy or Discovery Server; for example:

- A Discovering Node wants to discover given services related to a group of deployed sensors by sending a multicast to /.well-known/core. It gets responses for the sleeping sensors from the Proxy nodes.
- During commissioning phase, a discovering node queries a Discovery Server to find all the proxies providing a given service.

**REPORT Interaction:** On request of a Destination Node or because of configuration settings which have instructed the Node to do so, a Node sends a sequence of notifications of events to destination Node(s) (directly or via Proxy); for example:

- A battery-powered sensor sends a notification with "battery low" event directly to a designated Destination Node (Direct REPORT).

- A battery-powered occupancy sensor detects an event "people



present", switches on the radio and multicast a "ON" command to a group of lights (Direct REPORT).

- A battery-powered temperature sensor reports periodically the room temperature to a proxy Node (REPORT(A)). The proxy node reports to all associated HVAC destination nodes when the temperature change deviates from a predefined range (REPORT(B)).

WRITE Interaction: A node sends a request to a proxy to set a value.

- o A sleepy Node WRITES to the proxy; for example:
  - A battery-powered sensor wants to extend the registration lifetime of its delegated resource at the Proxy.
- o A configuring Node WRITES information to a Proxy; for example:
  - A configuring Node changes the reporting frequency of a deployed sensor by contacting the Proxy node to which the sensor is registered.
  - Sensor firmware is upgraded. A configuring Node pushes firmware data blocks to the Proxy, which pushes the blocks to the Sleepy Node.
  - A configuring Node adds a new subscription to an operational sensor via the Proxy. From that moment on, the new Node receives also the sensor events and status updates from the sensor.

READ Interaction: A node sends a read request to a node that returns a value.

- o Sleepy Node sends a read request to a server Node; for example:
  - A sensor (periodically) updates internal data tables by fetching it from a predetermined remote node.
  - A sensor (periodically) checks for new firmware with a remote node. If new firmware is found, the sensor switches to a non-sleepy operation mode, and fetches the data.
- o A sleepy node sends a read request to its proxy; for example:
  - A sensor (periodically) checks with his Proxy availability of configuration updates or changes of its delegated resources

(e.g. a sensor may detect in this way that a configuring Node has changed its name or modified its reporting frequency).

- o A reading Node sends a read request to a proxy; for example:
  - A Node (e.g. in the backend) requests the status of a deployed sensor, e.g. asking the sensor state and/or firmware version and/or battery status and/or its error log. The Proxy returns this information.
  - A Node requests a Proxy when a Sleepy sensor was 'last active' (i.e. identified as being awake) in the network.

## 2.2. Architecture

The architecture associated with the support of sleepy nodes is illustrated in Figure 2. Three High level interfaces are shown.




Figure 2: Architecture of sleepy node support

- o Direct interface: it allows the Sleepy Node to communicate directly to endpoints (i.e. for sending or reading information). The operations performed via this interface are always initiated by the Sleepy Node when its sleep period ends..
- o Delegate interface: via this interface the Proxy exposes the values of delegated resources to interested endpoints on behalf of the Sleepy Node. The same interface is used by endpoints which want to communicate with the Sleepy Node (e.g. for reading or writing information).
- o Synchronize interface: used by Sleepy Node and Proxy to synchronize values of delegated resources. Through this interface operations as discovery of the Proxy, registration, initialization and update of resources at the Proxy are performed, along with a de-registration operation to explicitly remove resources already registered to the Proxy.

The interfaces consist of a set of functions which together realize the interactions described in [Section 2.1](#).

Endpoints and the proxy communicate with a Resource Directory (RD) to discover resources of the sleepy node and delegated resources on the proxy (not shown in the Figure 2).

### [2.3](#). Example contents

The examples presented in this specification make use of a smart temperature sensor the resources of which are defined below using Link Format [[RFC6690](#)]. Three resources are dedicated to the Device Description (manufacturer, model, name) and one contains the current temperature in degree Celsius.

```
</dev/mfg >rt="ipso.dev.mfg";if="core.rp",  
</dev/mdl>rt="ipso.dev.mdl";if="core.rp",  
</dev/n>rt="ipso.dev.n";if="core.p",  
</sen/temp>rt="ucum.Cel";if="core.s"
```

## [3](#). Interactions involving Resource Directory

It is assumed that the Proxy has a resource type `rt="core.sp"`, where `sp` stands for sleepy proxy.

In order to become fully operational in a network and to communicate over the interfaces shown in Figure 2, a Sleepy Node and the Proxy need to perform operations via the Registration interface of the RD:

- Discovery of Proxy via RD. The sleepy node MAY discover the Proxy by sending a request to the RD to return all EP with `rt=core.sp`.
- Register existence of Proxy. When a RD is present and a Sleepy Node has registered itself to a Proxy (see [Section 4.2](#)), the Proxy MUST register the Sleepy Node at the RD and MUST keep this registration up-to-date.
- Register delegated resources. When a RD is present, the Proxy

MUST register the delegated resources at the RD and keep them up-to date.

A Configuring Endpoint (often part of a so-called Commissioning Tool) registers the services that are reported directly by the sleepy node in the resource directory, by registering the resource type and the multicast address. The multicast address can be associated with a group as described in [[I-D.ietf-core-resource-directory](#)].

A discovering Endpoint can discover one or more Sleepy Node resources via the Resource Directory.

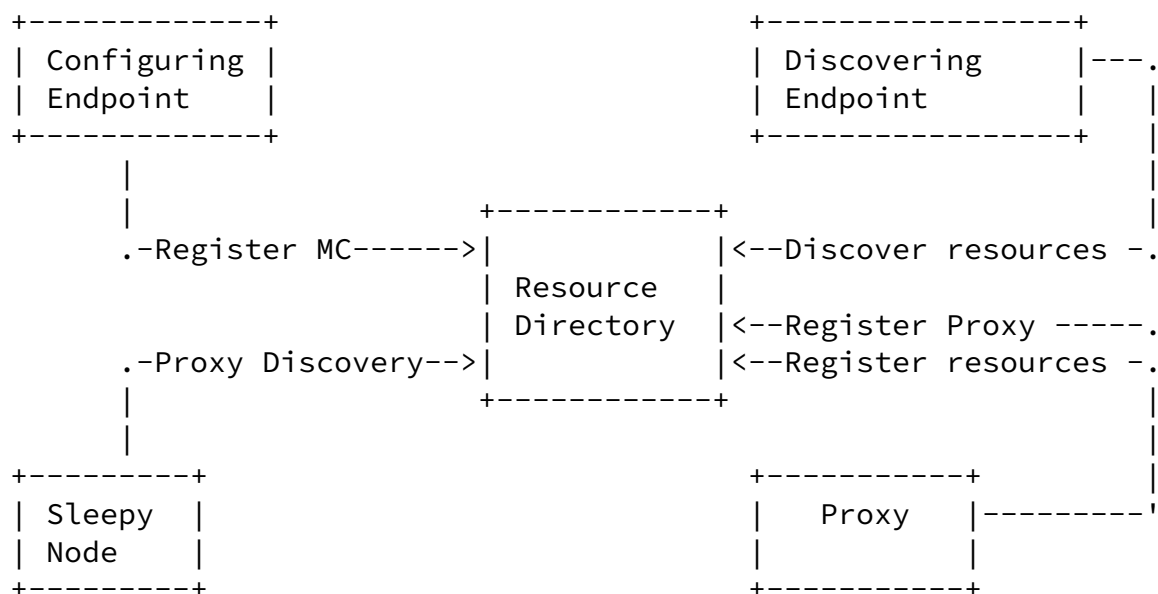


Figure 3: Interactions involving Resource Directory

#### [4.](#) Synchronize interface

The functions of the synchronize interface implemented by the Proxy are described in this section.

##### [4.1.](#) Sleepy node discovers proxy

A sleepy node can discover the proxy in two ways:

- via the CoAP interface [[RFC7390](#)] by sending a multicast message to discover an endpoint with rt=core.sp.
- via RD as already described in [Section 3](#).

The following example shows a sleeping endpoint discovering a proxy using this interface, thus learning that the base Proxy resource, where the sleepy node resources are registered, is at /sp.

Sleepy		Proxy
	----- GET /.well-known/core?rt=core.sp ----->	
	<----- 2.05 Content "</sp>; rt="core.sp" -----	

```
Req: GET coap://[ff02::1]/.well-known/core?rt=core.sp
Res: 2.05 Content
</sp>;rt="core.sp"
```

#### [4.2](#). Registration at a Proxy

Once a Sleepy Node has discovered a Proxy by means of one of the procedures described in [Section 4.1](#), the registration step can be performed. To perform registration, a Sleepy Node sends to the Proxy Node a CoAP POST request containing a description of the resources to be delegated to the Proxy as the message payload in the CoRE Link Format [[RFC6690](#)]. The description of the resource includes the Sleepy Node identifier, its domain and the lifetime of the registration.

Upon successful registration a Proxy creates a new delegated resource or updates an existing delegated resource and returns its location.

The resources specified by the Sleepy node during registration are created with path that has as prefix the base Proxy resource path (e.g. /sp). The registration interface MUST be implemented to be idempotent, so that registering twice with the same endpoint parameter does not create multiple delegated resources. The delegated resource SHOULD implement the Interface Type CoRE Link List defined in [[I-D.ietf-core-interfaces](#)]. A GET request on this resource MUST return the list of delegated resources for the corresponding sleepy node.

After successful registration, a Proxy SHOULD enable resource discovery for the new resources by updating its "/.well-known/core" resource. A Proxy MUST wait for the initial representation of a resource before it can be visible during resource discovery. The top level delegated resource MUST be published in "/.well-known/core" to enable the discovery of the resources via RD as described in [Section 3](#). Resources of a delegated container SHOULD be discoverable either directly in "/.well-known/core" or indirectly through gradual reveal from the delegated resource. The Web Link of a delegated resource MUST contain an "ep" attribute with the value of the End-Point parameter received during registration. If present, the End-Point Type parameter SHOULD also be mapped as a "rt" attribute.

A Proxy MAY be configured to register the Sleepy Node's resources in a RD. In this case, a Sleepy node MUST NOT register the resources in a RD by itself since it is the responsibility of the Proxy to perform the registration in the RD on behalf of the Sleepy Node. Since each Sleepy node may register resources with different lifetimes, a Proxy MUST register the resources of a given Sleepy Node in a dedicated path of the RD.

In case a Sleepy node delegates its own resources to more than one Proxy and each Proxy registers the Sleepy Node's resource in a RD, the RD entries from the different Proxies for the same Sleepy node risk to overlap.

To avoid this problem, a Proxy MUST create its own resource path to register the resources of a Sleepy node on the RD.

The new path name is typically formed by concatenating the Proxy's endpoint identifier with the path in use. This precaution ensures

that the ep identifier of a Sleepy node is unique for each resource path in the RD.

Implementation note: It is not recommended to reuse the value of the ep parameter in the URI of the delegated resource. This parameter may be a relatively long identifier to guarantee global uniqueness (e.g. EUI64) and would generate inefficient URIs on the Proxy where only a local handler is necessary.

The following example shows a sleepy node registering with a Proxy.

Sleepy	Proxy
--- POST /sp "</dev..." ----->	
<-- 2.01 Created Location: /sp/0 -----	

Req: POST coap://sp.example.org/sp?ep=0224e8fffe925dcf&rt=sensor

Etag: 0x3f

Payload:

```
</dev/mfg >;rt="ipso.dev.mfg";if="core.rp",  
</dev/mdl>;rt="ipso.dev.mdl";if="core.rp",  
</dev/n>;rt="ipso.dev.n";if="core.p",  
</sen/temp>;rt="ucum.Cel";if="core.s"
```

Res: 2.01 Created

Location: /sp/0

The delegated resource has been created with path /sp/0 on the Proxy in the example above. The path to the ep can be discovered as shown below:

Req: GET coap://sp.example.org/.well-known/core

Res: 2.05 Content

```
</sp>;rt="core.sp",  
</sp/0>;ep="0224e8fffe925dcf";rt="sensor"
```

A node can discover the delegated resources of the ep as shown below:

```

Req: GET coap://sp.example.org/sp/0
Res: 2.05 Content
Payload:
</sp/0/dev/mfg >;rt="ipso.dev.mfg";if="core.rp",
</sp/0/dev/mdl>;rt="ipso.dev.mdl";if="core.rp",
</sp/0/dev/n>;rt="ipso.dev.n";if="core.p",
</sp/0/sen/temp>;rt="ucum.Cel";if="core.s"

```

Once the resources are registered in the Proxy, the Proxy registers the delegated resources in the RD.

Proxy	RD
--- POST /rd "</sp/0..." ----->	
<-- 2.01 Created Location: /rd/6534 -----	

```

Req: POST coap://rd.example.org/rd?ep=0224e8fffe925dcf&
      rt=sensor

```

```

Etag: 0x6a
Payload:
</sp/0/dev/mfg >;rt="ipso.dev.mfg";if="core.rp",
</sp/0/dev/mdl>;rt="ipso.dev.mdl";if="core.rp",
</sp/0/dev/n>;rt="ipso.dev.n";if="core.p",
</sp/0/sen/temp>;rt="ucum.Cel";if="core.s"

```

```

Res: 2.01 Created
Location: /rd/6534

```

### [4.3.](#) De-registration at a Proxy

Sleepy node resources in the Proxy are kept active for the period indicated by the lifetime parameter. The Sleepy node is responsible for refreshing the delegated resource within this period using either the registration or update function (see [Section 4.5](#) of the

Synchronize interface. Once a delegated resource has expired, the Proxy deletes all resources associated to that resource and updates its `"/.well-known/core"` resource. When the Proxy resources are also registered in a RD, the RD and delegated resources are supposed to



have the same lifetime. Consequently, when the delegated resource expires, a Proxy MAY let the RD resource expire too instead of explicitly deleting it. When the delegated resource is deleted by means of explicit de-registration operation then also the RD resource MUST be explicitly removed.

A Proxy could lose or delete the delegated resource associated to a Sleepy node without sending an explicit notification (e.g. after reboot). A Sleepy node SHOULD be able to detect this situation by processing the response code while using the Sleepy node Operation or Update interface. Especially an error code "4.04 Not Found" SHOULD cause the Sleepy node to register again. A Sleepy node MAY also register with multiple proxies to alleviate the risk of interruption of service.

#### [4.4.](#) Initialization of delegated resource

Once registration has been successfully performed, the Sleepy Node must initialize the delegated resource. To send the initial contents (e.g. values, device name, manufacturer name) of the delegated resources to the Proxy, the Sleepy Node uses CoAP PUT repeatedly.

The basic interface is specified as follows:

Interaction: sleepy -> Proxy

Method: PUT

URI Template: `/[+location][+resource][?lt]`

URI Template Variables:

location := This is the Location path returned by the Proxy as a result of a successful registration.

resource := This is the relative path to a delegated resource managed by the registered sleepy node.

lt := Lifetime (optional). The number of seconds by which the lifetime of the whole delegated resource is extended. Range of 1-4294967295. If no lifetime is included, the current remaining lifetime stays unchanged.

Request Content-Type: Defined at registration

Response Content-Type: Defined at registration for GET method.  
application/link-format for PUT method if at least one of the  
mutable resources has been updated since the last PUT request.

Etag: The Etag option MAY be included to allow clients to validate a  
resource on multiple Proxies.

Success: 2.01 "Created", the request MUST include the initial  
representation of the delegated resource.

Success: 2.04 "Changed", the request MUST include the new  
representation of the delegated resource.

Success: 2.05 "Content", the response MUST include the current  
representation of the delegated resource.

Failure: 4.00 "Bad Request". Malformed request.

Failure: 5.03 "Service Unavailable". Service could not perform the  
operation.

The following example describes how a sleepy node can initialize the  
resource containing its manufacturer name just after registration.

Sleepy		Proxy
	--- PUT /sp/0/dev/mfg "acme" ----->	
	<-- 2.01 Created -----	

Req: PUT /sp/0/dev/mfg

Payload: acme

Res: 2.01 Created

The example below shows how a Sleepy node can indicate that it is  
supposed to send a temperature value at least every hour to keep its  
delegated resource active.

Internet-Draft

Sleepy Nodes

July 2015

```

Sleepy                                     Proxy
|                                         |
| --- PUT /sp/0/sen/temp?lt=3600 "22" -----> |
|                                         |
| <-- 2.04 Changed -----> |
|                                         |

```

Req: PUT /sp/0/sen/temp?lt=3600

Payload: 22

Res: 2.04 Changed

The use of repeated CoAP PUT can be avoided by writing all relevant resources into the Proxy in one operation by means of the Batch interface described in [[I-D.ietf-core-interfaces](#)]. After successful initialization, a Proxy SHOULD enable resource discovery for the new delegated resources by updating its /.well-known/core resource.

#### [4.5.](#) Sleepy Node updates delegated resource at Proxy

A Sleepy Node can update a delegated resource at the Proxy (REPORT A) using standard CoAP PUT requests on the delegated resource as shown in [Section 4.4](#).

When a Sleepy node sends a PUT request to update its resources, the response MAY contain a link-format payload. The payload does not directly relate to the target resource of the PUT request. Instead, it is a list of web links to resources that have been modified by clients since either the last PUT request or the last call to the modification check interface.

#### [4.6.](#) Sleepy Node READs resource updates from Proxy

This function allows a Sleepy Node to retrieve a list of delegated resources that have been modified at the Proxy by other nodes. The interface format for GET is the same as the one specified for PUT in [Section 4.4](#).

A configuring Node (EP) can update a resource in the Proxy. The sleepy node receives an indication of the changed resources as

specified in [Section 4.5](#).

The Sleepy Node can send GET requests to its Proxy on each delegated resource in order to receive their updated representation. The example in Figure 4 shows a configuration node which changes the name of a Sleepy Node at the Proxy. The Sleepy Node can then check and read the modification in its resource.

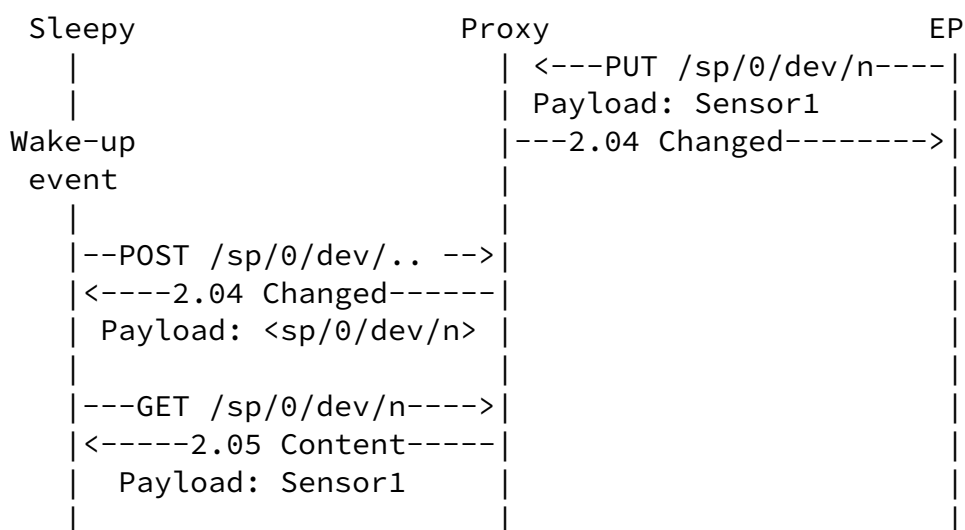


Figure 4: Example: A Sleepy Node READs resource updates from his Proxy

## 5. Delegate Interface

This section details the functions belonging to the delegate interface.

### 5.1. Discovering Endpoint discovers Sleepy Node at Proxy

Through this function, a Discovering Endpoint can discover one or more Sleepy Node(s) at a Proxy. In case a Resource Directory is not present, this is the only way to discover Sleepy Nodes. A CoAP client discovers resources owned by the Sleepy Node but hosted on the Proxy using typical mechanisms such as one or more GETs on the resource `/.well-known/core` [[RFC6690](#)].

Resource discovery between an Endpoint and a proxy or an Endpoint and

a RD needs special care to take into account the fact that resources from a sleeping node might appear duplicated. EPs SHOULD employ 2-step resource discovery by looking up sleeping nodes AND resource types to detect duplicate resources. EPs MAY use single-step resource discovery only if the Sleepy node can register with no more than one Proxy. An EP can use the "ep" link attribute as a filter on the "/.well-known/core" resource to retrieve a list of endpoints and detect duplicate sleeping nodes registered on multiple proxies. An EP can use the "ep" type of lookup to do the same on a RD. The result of endpoint discovery is then used to filter out duplicate resources returned from simple resource discovery.

The following example shows a client discovering the Sleepy nodes and learning that the sleepy node 0224e8fffe925dcf is registered on two Proxies.

EP	proxy1	proxy2
----- GET /.well-known/core?ep=* ----->	----->	
<----- 2.05 Content "</sp/0>..." -----		
<----- 2.05 Content "</sp/0>..." -----	-----	

```

Req: GET coap://[ff02::1]/.well-known/core?ep=*
Res: 2.05 Content
</sp/0>;ep="0224e8fffe925dcf"
Res: 2.05 Content
</sp/0>;ep="02004cffffe4f4f50"
</sp/1>;ep="0224e8fffe925dcf"

```

From the previous exchange and the next resource discovery request, the EP can infer that the resources coap://sp1/sp/0/sen/temp and coap://sp2/sp/1/sen/temp actually come from the same sleeping node with ep=0224e8fffe925dcf.

EP	proxy1	proxy2
- GET /.well-known/core?rt=ipso:ucum.Cel ->	----->	

	<----	2.05 Content "</sp/0>..."	-----		
	<----	2.05 Content "</sp/1>..."	-----	-----	

```

Req: GET coap://[ff02::1]/.well-known/core?rt=ucum.Cel
                                         &ep=0224e8fffe925dcf
Res: 2.05 Content
    </sp/0/sen/temp;rt="ucum.Cel"
Res: 2.05 Content
    </sp/1/sen/temp>;rt="ucum.Cel"

```

## 5.2. Proxy REPORTs events to Endpoint

This interface can be used by the Endpoint to receive event report message to Proxy (REPORT A) which further notifies it to interested Destination Endpoint(s) (REPORT B). This indirect reporting is useful for a scalable solution, e.g. there may be many interested subscribers but the Sleepy Node itself can only support a limited number of subscribers given its limits on battery energy. The mechanism according to which the Proxy forwards the event to Destination Endpoints (REPORT B) may be linked to a specific protocol (for example: publish/subscribe as in MQTT). A client interested in

the events related with a specific resource may send a CoAP GET to the Proxy, to obtain the last published state. If a Reading node is interested in receiving updates whenever the Sleepy Node reports new event to its Proxy, it can use observe at the Proxy for that specific resource.

A proxy using the CoAP protocol [[RFC7252](#)] SHOULD accept to establish a CoAP observation relationship between the delegated resource and a client as defined in [[I-D.ietf-core-observe](#)].

A Sleepy node may stop updating its delegated resources without explicitly removing its delegated resource (e.g. transition to another proxy after network unreachability detection). An Endpoint can detect this situation when the corresponding delegated resource has expired. Upon receipt of a response with error code 4.04 "Not Found", an Endpoint SHOULD restart resource discovery to determine if the resources are now delegated to another proxy.

The interface function is specified as follows:

Sleepy	Proxy	EP
	<- GET /sp/0/sen/temp - (observe)	
	-- 2.05 Content "22" ->	
- PUT /sp/0/sen/temp "23" ->		
<- 2.04 Changed -----		
	-- 2.05 Content "23" ->	

### 5.3. A Node WRITES to Sleepy Node via Proxy

A Configuring Node uses CoAP PUT to write information (such as configuration data) to the Proxy, where the information is destined for a Sleepy Node. Upon change of a delegated resource, an internal flag is set in the Proxy that the specific resource has changed. Next time the Sleepy Node wakes up, the Sleepy Node checks the Proxy for any modification of its delegated resources and reads those changed resources using CoAP GET requests, as shown in Figure 4. The allowed resources that a Configuring Node can write to, and the CoAP Content-Format of those CoAP resources, is determined in the initial registration phase.

The following example shows a commissioning tool (EP) changing the name of a Sleepy Node through a Proxy. The Sleepy Node detects this change right after updating its current temperature.

Sleepy	Proxy	EP
	<-- PUT /sp/0/dev/n ---	
	-- 2.04 Changed ----->	



- PUT /sp/0/sen/temp --->	
<- 2.04 Changed -----	
Payload: <sp/0/dev/n> ---	
- GET /sp/0/dev/n ----->	
<- 2.05 Content -----	

Req: PUT /sp/0/dev/n  
 Payload: "sensor-1"  
 Res: 2.04 Changed

Req: PUT /sp/0/sen/temp  
 Payload: "24"  
 Res: 2.04 Changed, Content-Type: application/link-format  
 Payload: "</sp/0/dev/n>"

Req: GET /sp/0/dev/n  
 Res: 2.05 Content  
 Payload: "sensor-1"

#### [5.4.](#) A Node READs information from Sleepy Node via Proxy

A Reading Node uses standard CoAP GET to read information of a Sleepy Node via a Proxy. However, not all information/resources from the Sleepy Node may be copied to the Proxy. In that case, the Reading Node cannot get direct access to resources that are not delegated to the Proxy. The strategy to follow in that case is to first WRITE to the Sleepy Node (via the Proxy, [Section 5.3](#)) a request for reporting this missing information; where the request can be fulfilled by the Sleepy Node the next time the Sleepy Node wakes up.

### [6.](#) Direct Interface

This section details the functions belonging to the direct interface.

#### [6.1.](#) Sleepy Node REPORTs events directly to Destination Node

When the Sleepy Node needs to report an event to Destination nodes or groups of Destination nodes present in the subscribers list, it

becomes Awake and then it can use standard CoAP POST unicast or multicast requests to report the event.

TODD: MC example

## 6.2. A Sleepy Node READs information from a Server Node

A Sleepy Node while Awake uses standard CoAP GET to read any information from a Server Node. While the Sleepy Node awaits a CoAP response containing the requested information, it remains awake. To increase battery life of Sleepy Nodes, such an operation should not be performed frequently.

## 7. Realization with PubSub server

The registration and discovery of the PubSub broker [[I-D.koster-core-coap-pubsub](#)] is covered to the same extent as discussed in this document. Not covered is the direct interaction between Sleepy node and destination nodes as described in [Section 6.1](#). The support from a server node to initialize resources or other information also represents an addition to PubSub broker.

In addition to the continuous updates provided by the PubSub broker, the ad-hoc query of values, the maintenance of operational parameters, the provision of direct update from Sleepy node to a node, the reliability aspects of the update, and the concept of groups are equally important topics that need consideration.

## 8. IANA Considerations

The new Resource Type (rt=) Link Target Attribute, 'core.sp' needs to be registered in the "Resource Type (rt=) Link Target Attribute Values" sub registry under the "Constrained RESTful Environments (CoRE) Parameters" registry.

## 9. Security Considerations

For the communication between Sleepy Node and proxy it MAY be sufficient to use Layer 2 (MAC) security without the recommended use of DTLS. However, it must be ascertained that the Sleepy Node can communicate only with a given secured Proxy. A Sleepy Node may obtain the Layer 2 network key using the bootstrapping mechanism described in [[I-D.kumar-6lo-selective-bootstrap](#)]. DTLS MUST be used over link-layer security for further transport-layer protection of messages between Regular Nodes and Proxies in the network. There are no special adaptations needed of the DTLS handshake to support Sleepy Nodes. During the whole handshake, Sleepy Nodes are required to remain awake to avoid that, in case of small retransmission timers,

Internet-Draft

Sleepy Nodes

July 2015

the other node may think the handshake message was lost and starts retransmitting. In view of this, the only key point, therefore, is that DTLS handshakes are not performed frequently to save on battery power. Based on the DTLS authentication, also an authorization method could be implemented so that only authorized nodes can e.g.

- Act as a Proxy for a Sleepy Node. (The Proxy shall be a trusted device given its important role of storing values of parameters for the delegated resources);
- READ data from Sleepy Nodes;
- WRITE data to Sleepy Nodes (via the Proxy);
- Receive REPORTs from Sleepy Nodes (direct or via Proxy).

## 10. Acknowledgements

Much of the text and examples in this document are copied from [[I-D.vial-core-mirror-server](#)]. Matthieu Vial has generously authorized us to use his text.

## 11. Changelog

RFC editor, please delete this section before publication.

From version 2 to version 3:

Introduced interfaces and copied examples and text from mirror server draft.

## 12. References

### 12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#), October 2010.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", [RFC 6690](#), August 2012.

[RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", [RFC 7252](#), June 2014.

Zotti, et al.

Expires January 7, 2016

[Page 24]

---

Internet-Draft

Sleepy Nodes

July 2015

[RFC7390] Rahman, A. and E. Dijk, "Group Communication for the Constrained Application Protocol (CoAP)", [RFC 7390](#), October 2014.

## [12.2](#). Informative References

[I-D.ietf-core-interfaces]

Shelby, Z. and M. Vial, "CoRE Interfaces", [draft-ietf-core-interfaces-02](#) (work in progress), November 2014.

[I-D.ietf-core-observe]

Hartke, K., "Observing Resources in CoAP", [draft-ietf-core-observe-16](#) (work in progress), December 2014.

[I-D.ietf-core-resource-directory]

Shelby, Z., Koster, M., Bormann, C., and P. Stok, "CoRE Resource Directory", [draft-ietf-core-resource-directory-03](#) (work in progress), June 2015.

[I-D.koster-core-coap-pubsub]

Koster, M., Keranen, A., and J. Jimenez, "Publish-Subscribe Broker for the Constrained Application Protocol (CoAP)", [draft-koster-core-coap-pubsub-01](#) (work in progress), March 2015.

[I-D.kumar-6lo-selective-bootstrap]

Kumar, S. and P. Stok, "Security Bootstrapping over IEEE 802.15.4 in selective order", [draft-kumar-6lo-selective-bootstrap-00](#) (work in progress), March 2015.

[I-D.vial-core-mirror-server]

Vial, M., "CoRE Mirror Server", [draft-vial-core-mirror-server-01](#) (work in progress), April 2013.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service

Discovery", [RFC 6763](#), February 2013.

#### Authors' Addresses

Teresa Zotti  
Philips Research  
High Tech Campus 34  
Eindhoven 5656 AE  
The Netherlands

Phone: +31 6 21175346  
Email: [teresa.zotti@philips.com](mailto:teresa.zotti@philips.com)

Zotti, et al.

Expires January 7, 2016

[Page 25]

---

Internet-Draft

Sleepy Nodes

July 2015

Peter van der Stok  
Consultant

Phone: +31 492474673  
Email: [consultancy@vanderstok.org](mailto:consultancy@vanderstok.org)

Esko Dijk  
Philips Research  
High Tech Campus 34  
Eindhoven 5656 AE  
The Netherlands

Phone: +31 6 55408986  
Email: [esko.dijk@philips.com](mailto:esko.dijk@philips.com)

