INTERNET-DRAFT
Intended Status: Proposed Standard
Expires: April 18, 2013

Mingui Zhang
Jie Dong
Huawei
Beichuan Zhang
The University of Arizona
October 15, 2012

Use Cases for Power-Aware Networks
draft-zhang-panet-use-cases-00.txt

Abstract

Power Aware NETwork (PANET) has attracted strong interest from both carriers and vendors. Several use cases are investigated in this document to exhibit the potential usage of PANET in both backbone and data center networks.

Table of Contents

1. Introduction

Networks are usually provisioned for peak hours and potential network failures and network devices are powered on all the time without consideration on energy efficient. In practice, however, the traffic load of a network is low most of the time and redundant network equipments are used for failure recovery occasionally.

In the past years, vendors had paid a great effort on improving the network energy efficiency at the device level: when the traffic load is low, a network equipment should accordingly operate with less power draw. However, network equipments have never become fully power proportional. Even few or no traffic is carried, a powered-on network device draws a considerable amount of power, which means energy is being wasted. There is an explicit gap for idle network devices to be shut down or put into sleeping state to save more energy. In order to fill this gap, the network control plane and management system should become power aware to coordinate network devices therefore the sleeping or off network devices do not bring disruption to the network.

This documents investigated several use cases on power aware network which include both backbone networks and data center networks. As for the energy efficiency of backbone networks, only intra-domain use cases are considered. Trying to be energy efficient in the inter-domain scale seems technically feasible, for now though, energy efficient solutions can easily end up lack of business motivation, this document leaves them for future study.

1.1. Conventions used in this document

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].

1.2. Terminology

PANET: Power Aware NETwork

2. Power Awareness in Backbone Networks

The IETF Energy Management (eman) Working Group works on the management of power-aware network devices. Basically, the power states of power-aware network devices are reported and recorded in MIB. However, there is a gap on how to make use of this kind of data to achieve energy efficient networks. With energy aware control plane [power-control], it becomes possible to make use of these measurements and power control ability to achieve the energy

efficiency of a whole network.

This section lists several use cases for backbone networks. Take a router system as an example, the start-up of it may take several minutes and the stabilization of it may takes much longer time. It is unrealistic to switch off and on a whole node in backbone networks frequently to achieve energy efficient, so this document only investigates the cases in which links (i.e., links' attached components) are shut-down for energy conservation.

2.1. Use Case 1: Sleeping Links

```
              +----------------+        +----------------+
              |Router          |        |          Router|
              |                |        |                |
              |                |        |                |
              |      +---------+        +---------+      |
              |      |Line-card|        |Line-card|      |
              |      |         |        |         |      |
              |      |     +---+  Link  +---+     |      |
              |      |     |NPU+--------+NPU|     |      |
              +------+-----+---+        +---+-----+------+
```

Figure 2.1: Power Aware Line-cards

The power draw on line-cards occupies a great portion in the total power consumption of a whole routing system. For high-end routers, this portion may be higher than 50%.

Network devices and their processing capacity are provisioned for worst cases such as traffic burst and busy hours. Most of the time, the network is lightly loaded. Unfortunately, the power consumption of network devices is not proportional to the traffic load on them. Even there is no load on them, there is still a considerable base power consumption. Unlike personal PCs which can be shut down or enter power saving modes (such as sleeping), network devices are powered on and running even they are idle. This reality means that the network is wasting powers.

The conception that "a link is put into sleep state" is frequently mentioned. In this document, this conception is formalized as follows. The coupled end-points (such as interfaces, NPU or whole line-cards) attached to a idle link (as shown in Figure 2.1) enter the sleeping mode to save energy.

Traffic aggregation are used to create the opportunity for more links to become idle. This process can be automated through the control plane, such as Traffic Engineering [GreenTE].

The essentials of this use case:

- o Devices to be Power Aware: Routers and their line-cards.

- o What actions to take: NMS measures the traffic load and power profile each link [eman]; Routers execute the green TE algorithm; Routers send out signals to trigger the power-on/power-off of a NPU on a line card.

2.2. Use Case 2: Composite Links

A composite link is logical link composed of multiple physical [I-D.ietf-rtgwg-cl-requirement]. The composite link attached end-points are responsible to map traffic onto the component links and maintain the state of the composite link. Power awareness can be applied to composite links as well. When the traffic volume on the composite link is low, some component links can be shut down to conserve energy consumption. When the traffic volume becomes high, the sleeping members links can be waken up to absorb the traffic load.

Compared to use case 1, the advantage of executing energy saving for composite link is that the connectivity of the composite link does not suffer unless all the component links are cut off. In this way, the control plane of the component link is not disrupted. In other words, when the end points of the composite link execute the energy conservation action, they can do it in a distributed way and decisions are made locally.

The essentials of this use case:

- o Devices to be Power Aware: Composite links attached end-points.

- o What actions to take: NMS measures the traffic load and power profile of component links; Attached end-points adaptively turn-on/turn-off component links according to the traffic load on the composite link.

Use case 1 and use case 2 may be combined in a real network to achieve more energy saving.

3. Power Aware in Data Center Networks

Servers, network devices (ICT equipments) are intensively placed in Data Centers. In comparison with ISP backbone networks, the operating of Data Center Networks are more power hungry. The growing amount of energy consumed by a Data Center has led to high operating costs.

Although non-ICT equipments, such as lighting and air conditioners,

in a Data Center consumes a notable large amount of energy as well, this section concentrate on talking about right sizing ICT equipments for energy conservation. Energy conservation of non-ICT equipments are out of the scope of this document.

3.1. Server Consolidation

With virtualization technology, Virtual Machines (VMs) can be consolidated to fewer physical servers while idled servers can be put into power saving mode or turned off to achieve energy conservation of the whole Data Center. Virtualization technology allows the administration of a Data Center Network respond rapidly to the fluctuating capacity requirements.

Through monitoring of the work load and power profile, the Data Center Network Management System (Orchestrator) can judge in which hours workload is high and in which hours workload is low. For example, nights are generally off-peak hours in which workload is at low level. Virtual machines can be moved to fewer servers therefore idle servers can powered off or put into sleep to save energy. Before peak hours (e.g., in the morning), sleeping or powered off servers should be waken up to accommodate more active virtual machines (VMs).

The essentials of this use case:

o Devices to be Power Aware: All servers in a data center.

o What actions to take: NMS measures the work load and power profile of servers; The orchestrator of a Data Center Network adaptively triggers the actions of VM migration, the power-off and power-on of servers according to the workload.

3.2. Power Aware Load Balancing Among Multiple Sites

An enterprise may have multiple data centers which spread out in different geographic locations. Generally, the ICT resources in these data centers are well replicated and a job can be directed to any of them for execution. These data centers form a large distributed Internet scale systems and the price of power supply for them varies between two different locations. The operating cost of such a system highly depends on the load balancing scheme. Being power aware, the system can map requests to locations where energy price is cheaper.

This use case makes use of the difference of the prices of power draw in different locations. The orchestration of data centers (the NMS) is responsible for monitoring the power profile and work load of the ICT devices located in different data centers.

The essentials of this use case:

- o Devices to be Power Aware: All ICT-equipments in a data center.

- o What actions to take: ICT devices report their work load and power consumption profile to NMS. The orchestration (NMS) of the Data Center Networks adaptively map the request onto sites in consideration of reducing the overall power bill of the system.

3.3. Elastic Network Infrastructure

Traffic load of a data center is generated by the work load on servers and applied on the network infrastructure. The changing work load determines that the traffic load varies as time goes on. However, network devices are always left on even though the traffic load fluctuates, which wastes energy inevitably when the traffic load is low.

Ideally, the network infrastructure is elastic and can fit the traffic pattern with minimum subset to minimize the energy consumption of the network infrastructure. For now, Data Center Networks generally work at layer 2. So this use case should be realized through manipulating switching paths, in comparison with the power aware routing at layer 3. Openflow switches of SDN may be utilized to achieve this goal [ElasticTree].

The essentials of this use case:

- o Devices to be Power Aware: All network equipments in a data center.

- o What actions to take: Network devices report their traffic load and power consumption profile to NMS. The orchestrator (NMS) of a Data Center Network adaptively build the switching paths upon the network infrastructure. The idled links are put into power saving mode (e.g., sleeping), so that the network infrastructure becomes energy efficient.

6. Security Considerations

This document raises no new security issues.

7. IANA Considerations

No new registry is requested to be assigned by IANA. RFC Editor: please remove this section before publication.

8. References

8.1. Normative References

[power-control] A. Retana, R. White, M. Paul, "A Framework and Requirements for Energy Aware Control Planes", draft-retana-rtgwg-eacp-00.txt, work in progress.

[eman] IETF, "Energy Management Working Group Charter", 2012, <http://datatracker.ietf.org/wg/eman/charter/>.

[I-D.ietf-rtgwg-cl-requirement] Villamizar, C., McDysan, D., Ning, S., Malis, A., and L. Yong, "Requirements for MPLS Over a Composite Link", draft-ietf-rtgwg-cl-requirement-07 (work in progress), June 2012.

8.2. Informative References

[GreenTE] Zhang, M. and et al. , "GreenTE: Power-Aware Traffic Engineering", ICNP 2010.

[Rate-Adaptation] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation," in Proceedings of USENIX NSDI, 2008.

[ElasticTree] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "ElasticTree: Saving Energy in Data Center Networks," in Proceedings of USENIX NSDI, 2010.

Author's Addresses

Mingui Zhang
Huawei Technologies Co.,Ltd
Huawei Building, No.156 Beiqing Rd.
Beijing 100095 P.R. China

Email: zhangmingui@huawei.com

Jie Dong
Huawei Technologies Co.,Ltd
Huawei Building, No.156 Beiqing Rd.
Beijing 100095 P.R. China

Email: jie.dong@huawei.com

Beichuan Zhang
The University of Arizona

Email: bzhang@cs.arizona.edu