

Individual Submission
Internet-Draft
Intended status: Informational
Expires: April 28, 2011

B. Ohlman
Ericsson
O. Strandberg
Nokia Siemens Networks
C. Dannewitz
University of Paderborn
A. Lindgren
SICS
R. Maglione
Telecom Italia
B. Ahlgren
SICS
D. Kutscher
NEC
October 25, 2010

Requirements for accessing data in network storage
draft-ohlman-decade-add-use-cases-reqs-02

Abstract

The DECOupled Application Data Enroute (DECADE) working group is specifying standardized interfaces for accessing in-network storage from applications to store, retrieve and manage data. The main objective is to provide a framework that is useful to P2P applications, without excluding other, possibly related applications that can benefit from accessing in-network storage. This memo presents Internet TV as a specific application scenario where access to in-network storage would be required and lists a set of concrete requirements that should be considered for the DECADE architecture and protocol specifications.

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](#)].

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 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](http://trustee.ietf.org/bcp78) 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
2.	Internet TV Scenario	5
2.1.	Detailed Scenario Description	5
2.2.	Summary	6
3.	Specific requirements	6
3.1.	Unique Naming of Information Objects	6
3.1.1.	Requirement	6
3.1.2.	Rationale	7
3.1.3.	Discussion	7
3.2.	Access to Information Objects	7
3.2.1.	Requirement	7
3.2.2.	Rationale	7
3.2.3.	Discussion	7
3.3.	Real-time Support	7
3.3.1.	Requirement	7
3.3.2.	Rationale	7
3.3.3.	Discussion	8
3.4.	Discovery service for DECADE in-network storage	8
3.4.1.	Requirement	8
3.4.2.	Rationale	8
3.4.3.	Discussion	8
3.5.	Multiple active DECADE Storage Servers	8
3.5.1.	Requirement	8
3.5.2.	Rationale	8
3.5.3.	Discussion	8
4.	IANA Considerations	9
5.	Security Considerations	9
6.	Acknowledgments	9
7.	Informative References	9
	Authors' Addresses	9

1. Introduction

The DECADE approach to access to in-network storage through standardized interfaces has been motivated by P2P application scenarios where it can be beneficial to refer peers to in-network storage system in a convenient network-topological location in order to enhance data exchange for a given P2P application.

One specific example would be a P2P application instance in a home network connected via an ADSL access network. In a traditional P2P approach, this application instance would download chunks from other peers and serve chunks to other peers in parallel. With an asymmetric uplink, the application instance (and other hosts on the home network) are likely to experience uplink congestion, if the application instance is selected by a substantial number of peers and has to serve data to them. In situations with prevalent asymmetric access links, the P2P session (all peer application instances) would also be limited in the achievable downlink speed because the aggregate uplink bandwidth is not sufficient.

DECADE in-network storage servers are expected to be helpful in such scenarios, because peers could be referred to appropriate storage servers for downloading some content, thereby offloading traffic from the capacity-limited access uplinks.

Such storage servers are expected to provide interfaces for storing, retrieving, and managing data. The general concept is that, in a distributed P2P application, some instance would upload data to a storage server and then refer other instances, such as P2P peers, to that data, for instance by passing a certain URI. In addition, there could be interactions for allocating capacity on servers for certain applications, for deleting data etc.

This document argues that, while such a system would be very useful for P2P applications, there are others, related applications that could benefit from in-network storage in just the same way. Though only applications that does not lead to a completely new set of requirements should be taken into account. This document argues that it should be possible to extend the DECADE work to include certain other applications without increasing the overall complexity of the solution.

Specifically, this document describes the in-network-storage aspects of Internet TV -- an important application today, that can significantly benefit from in-network-caching. We argue that it seems negligent to exclude Internet TV from leveraging DECADE in-network-storage and describe specific use cases for accessing in-network-storage in such a scenario. Moreover, we present a set of

requirements that should be met by the DECADE architecture design in order guarantee its applicability to such applications. We propose that these requirements be added to the DECADE requirements specification.

In this memo, we align the requirement description to the layout used in [[I-D.gu-decade-reqs](#)]. [Section 2](#) describes the Internet TV scenario, and [Section 3](#) presents corresponding requirements that should be considered to ensure a broad enough applicability of the DECADE framework.

[2. Internet TV Scenario](#)

Internet TV is a general term to refer to different kinds of systems or applications where video is delivered to (mostly) home network devices for immediate rendering or storing. In this memo, we refer to the distribution of video content, mainly focusing on Video-on-Demand (VoD) services and user-generated content.

VoD services are commonly widespread in many service providers' networks. This scenario is characterized by the need to support an efficient large-scale distribution of video, possibly with a fairly high degree of replicated contents, to a multiplicity of fixed and mobile users. By supporting this application with DECADE protocols, video content can be retrieved from the in-network storage, achieving a number of benefits. The originating servers can be relieved from most of the load, since popular content will be automatically available in the in-network storage, closer to the users. Improved network efficiency will be achieved, reducing the traffic load in the upstream network segments. Moreover user experience, also for mobile users, can be improved.

[2.1. Detailed Scenario Description](#)

A well-known issue with Internet TV applications such as YouTube is the flash crowd problem. That is also an example of a problem which could be significantly eased if in-network storage is used to provide users with locally available copies rather than all requesting the data from the source. This can be extra beneficial for services with real-time (or near real-time) components as traditional pre-caching solutions can be difficult to use then.

A particular interesting Internet TV variant is "hybrid Internet TV" based on an Internet TV distribution service that is a hybrid between traditional CDN and a P2P service. Such a service would distribute content from central servers, make use of CDN caches on the way and finally use the end hosts/STB as caches for the P2P part of the

application.

If only the P2P application in the host/STB can store content in the DECADE storage, the content first has to be downloaded from the Internet TV server/CDN cache over the access link to the host/STB and then uploaded, over the same access link, to the DECADE storage before any peer in the P2P part of the application can access it from the DECADE storage (instead of downloading it from the client).

To avoid this, it should be possible for the DECADE storage to 'cache' the content when the first download to the host/STB is done. That would mean the content never have to travel over the capacity-limited uplink. For this to be feasible, one requirement is that the Internet TV service can prompt the DECADE storage that certain content should be cached on its way to the host. Having such functionality, that allows a host to get content from the DECADE storage of neighboring host rather than from a central server, would of course also offload the core network in the same way a traditional CDN does.

2.2. Summary

The DECADE architecture and protocol specifications should take the hybrid Internet TV scenario into account to ensure a reasonable level of generality of the DECADE in-network storage. While P2P-specific requirements should be considered, DECADE should not be unnecessarily limited to it.

Specifically, dissemination applications of streaming type (some with real-time or close to real-time requirements) should be supported by DECADE as they can cause significant load on the network. The network load could be reduced significantly for these types of applications if copies stored locally in the network could be used instead of always fetching data from the source.

3. Specific requirements

3.1. Unique Naming of Information Objects

3.1.1. Requirement

When a DECADE client in a certain application context stores an information object in DECADE storage servers, the object **MUST** be addressable by a unique name across different application contexts.

3.1.2. Rationale

There is a need for unique naming to enable different application instances to refer to information objects using a name (that may have been provided to them by another DECADE client). Such unique naming is essential for efficient cache handling and can serve for de-duplication.

3.1.3. Discussion

Unique naming can be achieved in different ways. Names can assigned from some (structured) names, for instance by URIs. Names can also be generated, for instance by calculating hashes of the object's content.

The detailed syntax and semantics of DECADE names (and the actual standardization requirements) are for further study.

3.2. Access to Information Objects

3.2.1. Requirement

It MUST be possible to access data stored on DECADE storage servers as complete information objects, such as a named video file.

3.2.2. Rationale

In a video-on-demand caching use case, the client application should be enabled to retrieve the complete object in one transaction and should not be required to download individual chunks.

3.2.3. Discussion

This does not necessarily impose implications on the way that the storage servers stores the object.

3.3. Real-time Support

3.3.1. Requirement

The DECADE storage service MUST support real-time applications in a way that a resource that is being uploaded is already available for download.

3.3.2. Rationale

For larger objects or chunks, it is not acceptable if a DECADE client has to upload the complete resource first, before other clients can

start downloading it.

3.3.3. Discussion

This requirement should also be important for P2P live streaming.

3.4. Discovery service for DECADE in-network storage

3.4.1. Requirement

When a DECADE client attach to a DECADE enabled network there SHOULD be a discovery service that can tell a DECADE client where in-network storage servers can be found.

3.4.2. Rationale

To minimize manual configuration of the DECADE clients, a discovery service, similar to DHCP , should be provided in the DECADE enabled network.

3.4.3. Discussion

In particular, this simplifies the administration of the DECADE in-network storage for a user that roams to a visited network.

3.5. Multiple active DECADE Storage Servers

3.5.1. Requirement

A DECADE client SHOULD be able to use multiple in-network storage servers at the same time.

3.5.2. Rationale

One example of when this is needed is when a user/client roams to another network, then it is reasonable to assume that the currently used in-network storage remains active for a certain time not to disrupt ongoing communication sessions at the same time as another in-network storage might immediately be needed in the new network.

3.5.3. Discussion

A user of DECADE in-network storage who roams to a visited network could potentially cause very inefficient access to that user's DECADE storage. It is therefore essential that the user is able to acquire new DECADE storage which is better located in the visited network. Usage that could result in such inefficiencies is communication with other users locally in the same network, for example as part of a

small meeting or large event (fair, sports event, etc).

A related issue is the possibility to migrate content from one DECADE storage to another when roaming. We believe that this is covered by the requirements on Efficient Transfer ([section 3.3](#)) and Communication among In-network Storage Elements ([section 4.3](#)) of [[I-D.gu-decade-reqs](#)].

[4.](#) IANA Considerations

This document has no requests to IANA.

[5.](#) Security Considerations

The re-use of copies in the network part of DECADE will require that appropriate access control mechanisms are designed.

[6.](#) Acknowledgments

We would like to thank all persons participating in the Network of Information work package in the EU FP7 projects 4WARD and SAIL for contributions and feedback to this document.

[7.](#) Informative References

[I-D.gu-decade-reqs]

Yingjie, G., Bryan, D., Yang, Y., and R. Alimi, "DECADE Requirements", [draft-gu-decade-reqs-05](#) (work in progress), July 2010.

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

Authors' Addresses

Borje Ohlman
Ericsson
Stockholm
Sweden

Email: Borje.Ohlman@ericsson.com

Ove Strandberg
Nokia Siemens Networks
Linnoitustie 6
Espoo
Finland

Email: ove.strandberg@nsn.com

Christian Dannewitz
University of Paderborn
Paderborn
Germany

Email: cdannewitz@upb.de

Anders Lindgren
Swedish Institute of Computer Science
Stockholm
Sweden

Email: andersl@sics.se

Roberta Maglione
Telecom Italia
Turin
Italy

Email: roberta.maglione@telecomitalia.it

Bengt Ahlgren
Swedish Institute of Computer Science
Stockholm
Sweden

Email: bengta@sics.se

Dirk Kutscher
NEC Laboratories Europe
Kurfuersten-Anlage 36
Heidelberg
Germany

Email: kutscher@neclab.eu