DECADE
Internet-Draft
Intended status: Informational
Expires: February 14, 2013

Y. Gu
Huawei
D. Bryan
Ethernot.org
Y. Yang
Yale University
P. Zhang
Tsinghua University/Yale University
R. Alimi
Google

August 13, 2012

DECADE Requirements
draft-ietf-decade-reqs-08

Abstract

The target of the DECoupled Application Data Enroute (DECADE) system is to provide an open and standard in-network storage system for applications, primarily P2P (peer-to-peer) applications, to store, retrieve and manage their data. This draft enumerates and explains requirements, not only for storage and retrieval, but also for data management, access control and resource control, that should be considered during the design and implementation of a DECADE-compatible system. These are requirements on the entire system; some of the requirements may eventually be implemented by an existing protocol with/without some extensions (e.g., a protocol used to read and write data from the storage system). The requirements in this document are intended to ensure that a DECADE-compatible system architecture includes all of the desired functionality for intended applications.

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 February 14, 2013.

Copyright Notice

Table of Contents

1. Introduction 5
2. Terminology 6
2.1. DECADE-compatible Client 6
2.2. DECADE-compatible Server 6
2.3. DECADE Storage Provider 6
2.4. DECADE Account 6
2.5. Resource Provider 6
2.6. Resource Consumer 6
2.7. Content Distribution Application 6
2.8. Target Application 7
2.9. Application End-Point 7
2.10. Data Object 7
2.11. DECADE-compatible System 7
2.12. DECADE Resource Protocol (DRP) Functions 7
2.13. DECADE Standard Data Transfer Protocol (SDT) Functions . . 7
3. System and Protocol Components 8
4. Requirements Structure 9
5. Data Object Requirements 9
5.1. Data Name Uniqueness 9
5.2. Verifiable Name-Object Binding 10
5.3. Data Object Size 10
5.4. Data Object Attributes 10
5.5. Application-defined Object Properties and Metadata 11
6. Access and Authorization Requirements 11
6.1. Provider Access 11
6.2. Secure Authorization 11
6.3. Consumer Access 12
6.4. Provider Authorization When Offline 12
6.5. Local Authorization 12
6.6. Access Control Granularity 12
6.7. Default Access Permissions 12
6.8. Connectivity Supporting NAT and Firewall Traversal 13
6.9. DECADE Server Discovery 13
7. Data Transfer Requirements 13
7.1. Negotiable Standard Data Transport Protocol 13
7.2. Atomic or Partial Read/Write 14
7.3. Secure Data Transport 14
7.4. Concurrent Read 14
7.5. Concurrent Write 14
7.6. Read Before Write Complete 15
7.7. Redirection of Transport 15
8. Resource Control Requirements 16
8.1. Multiple Applications Sharing Resources 16
8.2. Per-Client Resource Policy 16
8.3. Distributed Resource Sharing Specification 16
8.4. Resource Set 17

9. Error and Failure Handling Requirements 17
9.1. Illegal Data Object 17
9.2. Invalid Access Authorization 18
9.3. Insufficient Resources 18
9.4. Overload Condition 18
9.5. Attack Mitigation 19
10. Management Info Requirements 19
10.1. Account Status 19
11. Security Considerations 19
11.1. Authentication and Authorization 20
11.2. Confidentiality 20
11.3. Attack Mitigation 20
12. IANA Considerations 20
13. References 20
13.1. Normative References 20
13.2. Informative References 20
Appendix A. Acknowledgments 21
Authors' Addresses 21

1. Introduction

The object of the DECoupled Application Data Enroute (DECADE) system is to provide an open and standard in-network storage for content distribution applications, where data is typically broken into one or more chunks and then distributed. This may already include many types of applications including P2P applications, IPTV (Internet Protocol Television), and VoD (Video on Demand). (For a precise definition of the applications targeted in DECADE system, see the definition for Target Application in Section 2.) Instead of always transferring data directly from a source/owner client to a requesting client, the source/owner client can write to and manage its content on its in-network storage. The requesting client can get the address of the in-network storage pertaining to the source/owner client and read data from the storage.

This draft enumerates and explains the rationale behind specific requirements on the protocol design and on any data store implementation that may be used to implement DECADE servers that should be considered during the design and implementation of a DECADE-compatible system. As such, it does not include general guiding principles. General design considerations, explanation of the problem being addressed, and enumeration of the types of applications to which a DECADE-compatible system may be suited is not considered in this document. For general information, please see [RFC6646] and [I-D.ietf-decade-arch].

This document enumerates the requirements to enable target applications to utilize in-network storage. In this context, using storage resources includes not only basic capabilities such as writing, reading, and managing data, but also controlling access for particular remote clients with which it is sharing data. Additionally, we also consider controlling the resources used by remote clients when they access data as an integral part of utilizing the network storage.

This document discusses requirements pertaining to DECADE-compatible protocol(s). In certain deployments, several logical in-network storage systems could be deployed (e.g., within the same administrative domain). These in-network storage systems can communicate and transfer data through internal or non-standard communication messages that are outside of the scope of these requirements, but they should use DECADE-compatible protocol(s) when communicating with other DECADE-compatible in-network storage systems.

2. Terminology

This document uses the term 'In-network storage' which is defined in [RFC6646].

This document also defines these additional terms:

2.1. DECADE-compatible Client

A DECADE-compatible client uploads and/or retrieves data from DECADE-compatible servers. We use the shorter term "client" if there is no ambiguity.

2.2. DECADE-compatible Server

A DECADE-compatible server stores data inside the network, and thereafter manages both the stored data and access to those data. We use the shorter term "server" if there is no ambiguity.

2.3. DECADE Storage Provider

A DECADE Storage Provider, or Storage Provider for short, deploys and/or manages DECADE-compatible server(s) within a network.

2.4. DECADE Account

An account of a DECADE-compatible server has associated cryptographic credentials for access control, and resource allocation attributes on the server.

2.5. Resource Provider

A client which has the account cryptographic credentials of a DECADE account at a DECADE-compatible server. We use the short term "Provider" if there is no ambiguity.

2.6. Resource Consumer

A client which tries to access a DECADE account but does not have the account's cryptographic credentials. We use the short term "Consumer" if there is no ambiguity.

2.7. Content Distribution Application

A Content Distribution Application is an application (e.g., P2P, traditional CDN, or hybrid P2P/CDN) designed for dissemination of a large amounts of content (e.g., files, or video streams) to multiple content consumers. Content Distribution Applications may divide

content into smaller blocks for dissemination.

2.8. Target Application

An application with that includes a DECADE-compatible client along with other application functionalities to explicitly control the usage of resources of DECADE-compatible servers to deliver content to other users. A primary type of Target Application we consider is Content Distribution Applications. A Target Application divides content into smaller blocks for more flexible distribution (e.g., over multiple application-level paths). We use the term Target Application to refer to the type of applications that are explicitly (but not exclusively) supported by DECADE system.

2.9. Application End-Point

An Application End-Point is an instance of a Target Application. A particular Application End-Point might be a Provider, a Consumer, or both. For example, an Application End-Point might be an instance of a video streaming client, or it might be the source providing the video to a set of clients.

2.10. Data Object

A data object is the unit of data stored and retrieved from a DECADE-compatible server. The data object is a string of raw bytes. The server maintains metadata associated with each data object, but the metadata is not included in the data object.

2.11. DECADE-compatible System

A system which is composed of DECADE-compatible clients, DECADE-compatible servers and in-network storage. A DECADE-compatible system MUST obey all the requirements defined in this document.

2.12. DECADE Resource Protocol (DRP) Functions

A set of functions for communication of access control and resource scheduling policies from a DECADE client to a server, as well as between servers.

2.13. DECADE Standard Data Transfer Protocol (SDT) Functions

A set of functions to be used to transfer data objects to and from a DECADE server.

3. System and Protocol Components

To organize requirements, we consider that a DECADE-compatible system consists of two logical sets of functions, as shown in Figure 1. The first set of functions, which we refer to as the DECADE Resource Protocol (DRP) functions, is responsible for communication of access control and resource scheduling policies from a client to a server, as well as between servers. A DECADE-compatible system will include exactly one DRP for interoperability and a common format through which these policies can be communicated.

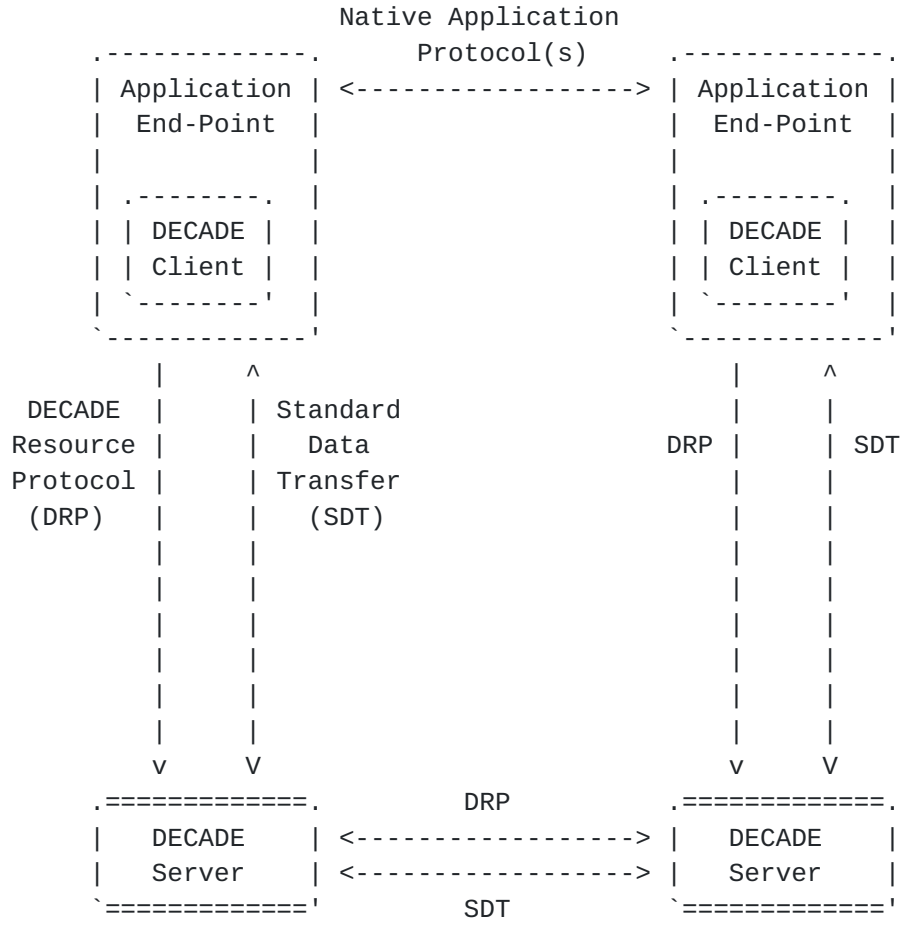

Figure 1: Protocol Components and Generic Flow

Second, the second set of functions, referred to as the Standard Data Transfer (SDT) functions, will be used to transfer data objects to and from a server. A DECADE-compatible system may support multiple SDT's. If there are multiple SDT's, a negotiation mechanism will be used to determine a suitable SDT between the client and server.

The two sets of functions (DRP and SDT) will be either separate or combined on the wire. If they are combined, DRP messages can be

piggy-backed within some extension fields provided by certain SDT protocols. In such a scenario, DRP is technically a data structure (transported by other protocols), but it can still be considered as a logical protocol that provides the services of configuring DECADE-compatible resource usage. If the protocols are physically separate on the wire, DRP can take the form of a separate control connection open between the a DECADE-compatible client and server. Hence, this document considers SDT and DRP as two separate, logical functional components for clarity.

4. Requirements Structure

This document specifies the requirements for the DECADE DRP and SDT functions, either existing ones or new ones, and storage system to enable Target Applications to make use of storage within the network, leaving specific storage system considerations to the implementation of the storage servers as much as possible. For this reason, we consider two primary categories of requirements:

- Protocol Requirements: Protocol requirements for Target Applications to make use of in-network storage within their own data dissemination schemes. Development of these requirements is guided by a study of data access, search and management capabilities used by Target Applications. These requirements may be met by a combination of existing protocols and new protocols.

- Storage Requirements: Functional requirements necessary for the back-end storage system employed by the DECADE server. Development of these requirements is guided by a study of the data access patterns used by Target Applications. These requirements should be met by the underlying data transport used by DECADE system. In this document, we use "data transport" to refer to a protocol used to read and write data from in-network storage.

This specification discusses the requirements of functionality implemented with a storage system and within applications, to permit interoperable communications concerning the manipulation of stored content.

5. Data Object Requirements

5.1. Data Name Uniqueness

REQUIREMENT(S): Each Data Object should be named to allow access. DECADE-compatible protocol(s) MUST support a data object naming scheme that ensures a high probability of uniqueness, with no coordination among multiple Storage Providers. In other words, two Data Objects with the same name should be the same content with high probability. A DECADE-compatible server SHOULD be able to respond to a DECADE-compatible client with an error indicating potential name conflicts.

RATIONALE: Although the intention of unique names is to avoid name collisions, it does not have to be an absolutely zero possibility. Hence, it is required to provide a collision handling mechanism.

EXCEPTION: While a DECADE-compatible server is overloaded or consider a request as an attack, it does not to generate a response to indicate name collisions.

5.2. Verifiable Name-Object Binding

5.3. Data Object Size

REQUIREMENT(S): DECADE MUST allow for efficient storage and data transfer of small data objects (e.g., 16KB) without large control overhead.

RATIONALE: Though Target Applications are frequently used to share large amounts of data (e.g., continuous streams or large files), the data itself is typically subdivided into smaller data objects (chunks) for flexibility (e.g., reliability and multi-path transmission).

5.4. Data Object Attributes

REQUIREMENT(S): DECADE MUST support the ability to associate a set of system attributes with a data object with a scope local to a DECADE-compatible server. In particular, the set MUST include time-to-live (or expiration time), creation timestamp, object size, and object type. A DECADE-compatible client, with access permission, MUST be able to query the set of system attributes. The transmission of the attributes MUST use an operating system-independent and architecture-independent standard format. An ability to extend the set of attributes MUST exist.

RATIONALE: The values of attributes associated with a data object are local to a particular DECADE-compatible server. These attributes may be used as hints to the storage system, internal optimizations, or as additional information query-able by DECADE-

compatible clients. The particular requirement for including time-to-live (TTL) is that a data object written by a DECADE-compatible client may be usable only within a certain window of time, such as in a live-streaming P2P application. Providing a time-to-live value for a data object (e.g., at the time it is written) can reduce management overhead by avoiding many 'delete' commands sent to DECADE-compatible server. The server may still retain a data object for bandwidth optimization, but this should be guided by the privacy policy of the DECADE Storage Provider.

5.5. Application-defined Object Properties and Metadata

REQUIREMENT(S): DECADE-compatible clients and DECADE-compatible servers MUST NOT be able to associate Application-defined properties (metadata) with data objects beyond what is provided by Section 5.4.

RATIONALE: Associating key-value pairs that are defined by Target Applications with data objects introduces substantial complexity. If Target Applications wish to associate additional metadata with a data object, possible alternatives include (1) managing such metadata within the Target Application itself, (2) storing metadata inside the data object, or (3) storing metadata in a different data object at the DECADE-compatible server.

6. Access and Authorization Requirements

6.1. Provider Access

REQUIREMENT(S): A Provider MUST be able to access the resources of its account.

RATIONALE: After a DECADE-compatible client owns an account on a DECADE-compatible server, it should be able to read data from and write data to the server.

6.2. Secure Authorization

REQUIREMENT(S): Access to an account on a server MUST be granted to a provider based on cryptographic security.

RATIONALE: DECADE-compatible clients may be operating on hosts without constant network connectivity or without a permanent attachment address (e.g., mobile devices). To support access control with such hosts, DECADE-compatible servers must support access control policies that use cryptographic credentials, not simply by tying access to IP addresses.

6.3. Consumer Access

REQUIREMENT(S): A Provider MUST be able to indicate to its server on whether a Consumer can access its subscribed resources.

RATIONALE: Endpoints in Target Applications may choose different servers. Thus, to be useful by Target Applications, a DECADE-compatible client must be able to specify policies on whether other DECADE-compatible clients can access its resources. The other clients may or may not be known to the server.

6.4. Provider Authorization When Offline

REQUIREMENT(S): A Provider MUST be able to grant access to a Consumer even if the Provider is not actively running or connected to its DECADE-compatible server.

RATIONALE: If an application desires, it can authorize other clients to access its storage even after the application exits or network connectivity is lost. An example use case is mobile scenarios, where a client can lose and regain network connectivity often.

6.5. Local Authorization

REQUIREMENT(S): A Provider MUST be able to indicate, without contacting its server, access control policies for Consumers. DECADE-compatible server MUST be able to authenticate the access control policies in this situation.

RATIONALE: This requirement is related with the preceding requirement, but in a perspective (i.e., protocol design). See discussions in Section 8.3.

6.6. Access Control Granularity

REQUIREMENT(S): A Provider MUST be able to control which individual clients are authorized to read/write which particular data objects from/to its in-network storage.

RATIONALE: A Target Application should able to conduct access control on the granularity of individual clients, individual data objects.

6.7. Default Access Permissions

REQUIREMENT(S): Unless read or write access is granted by a Provider, the default permission MUST be no access.

RATIONALE: This requirement is to protect client privacy by default.

6.8. Connectivity Supporting NAT and Firewall Traversal

REQUIREMENT(S): A client that is authorized to access a server MUST be supported to conduct NAT (Network Address Translation) and firewall traversal. In particular, network connections between a DECADE-compatible client and a DECADE-compatible server MUST be initiated by the client (i.e., the server must not initiate a connection to the client).

RATIONALE: Firewalls and NATs are widely used in the Internet today, both in ISP (Internet Service Provider) and Enterprise networks and by consumers. Many firewalls and NATs are configured by default to block incoming connections, which helps to mitigate security risks. Deployment of a DECADE-compatible system must not require manual modifications to such devices. This requirement applies to both potential new protocol that may be developed by the DECADE Working Group and any data transport used with DECADE protocol.

6.9. DECADE Server Discovery

REQUIREMENT(S): A mechanism for a Provider to discover and connect to its assigned server MUST be supported. The discovery SHOULD leverage existing mechanisms and protocols wherever possible. No new discovery mechanism will be defined unless there is enough evidence that no existing mechanism can work.

RATIONALE: Existing protocols such as DNS and DHCP are widespread. Using existing protocols, or combinations of the protocols that have been specified in other contexts, is strictly preferred over developing a new discovery protocol.

7. Data Transfer Requirements

7.1. Negotiable Standard Data Transport Protocol

REQUIREMENT(S): A DECADE-compatible client MUST be able to negotiate with a DECADE-compatible server about which protocol it can use to read data from and write data. DECADE MUST specify at least one mandatory transport protocol to be supported by implementations; usage of a different protocol may be selected via negotiation.

RATIONALE: Since typical data transport protocols (e.g., NFS and WebDAV) already provide read and write operations for network storage, it may not be necessary to define such operations in a new DECADE protocol. However, because of the particular application requirements and deployment considerations, different applications may support different protocols. Thus, a DECADE client must be able to select an appropriate protocol also supported by the in-network storage.

7.2. Atomic or Partial Read/Write

REQUIREMENT(S): A DECADE-compatible server MUST support the ability to read/write a complete data object in one request. It MAY support range read/write.

RATIONALE: Depending on the object size (e.g., chunk size) used by a Target Application, the application may need to send data to the DECADE-compatible server in multiple round.

7.3. Secure Data Transport

REQUIREMENT(S): A secure transport MUST be implemented for all communications between a DECADE-compatible client and DECADE-compatible server.

RATIONALE: Target Applications may wish to write sensitive data. To satisfy this use case, the communication between a DECADE-compatible client and DECADE-compatible server must be transported over a secure transport protocol (e.g., SSL/TLS).

7.4. Concurrent Read

REQUIREMENT(S): A DECADE-compatible server MUST allow for multiple simultaneous readers for a data object.

RATIONALE: One characteristic of Target Applications is the ability to upload an object to multiple clients. Thus, it is natural for DECADE-compatible server to allow multiple readers to read the same object concurrently.

7.5. Concurrent Write

REQUIREMENT(S): A DECADE-compatible server MUST NOT allow multiple simultaneous writers for the same object. A DECADE-compatible server SHOULD respond to each of the writers with an error.

RATIONALE: This avoids data corruption in a simple way while remaining efficient. Alternately, the DECADE-compatible server would need to either manage locking, handle write collisions, or merge data, all of which reduce efficiency and increase complexity.

EXCEPTION: While a DECADE-compatible server is overloaded or considers a request as an attack, it does not generate a response.

7.6. Read Before Write Complete

REQUIREMENT(S): A DECADE-compatible server MAY allow readers to read a data object before it has been completely written. In case of a write error in such a case, the DECADE-compatible server SHOULD respond to the reading client with an error indicating that the write has failed.

RATIONALE: Some Target Applications (in particular, P2P streaming) can be sensitive to latency. A technique to reduce latency is to remove store-and-forward delays for data objects (e.g., make the object available before it is completely written). Appropriate handling for error conditions (e.g., a disappearing writer) needs to be specified.

EXCEPTION: While a DECADE-compatible server is overloaded or considers a request as an attack, it does not generate a response.

7.7. Redirection of Transport

REQUIREMENT(S): A DECADE-compatible server SHOULD be able to redirect requests to another DECADE-compatible server. This may either be in response to an error, failure, overload, or to support capabilities such as load balancing.

RATIONALE: A DECADE-compatible server may opt to redirect requests to another server to support capabilities such as load balancing, or if the implementation decides that another DECADE-compatible server is in a better position to handle the request due to either its location in the network, server status, or other consideration.

EXCEPTION: A DECADE-compatible server can be configured by its service provider to support or not support redirection functionality.

8. Resource Control Requirements

8.1. Multiple Applications Sharing Resources

REQUIREMENT(S): A client MUST be able to indicate to a DECADE-compatible server about resource sharing policies among multiple Target Applications being run/managed by the same client.

RATIONALE: A client owning a DECADE account may provide the account's cryptographic credentials to multiple Providers of multiple target applications. For example, the client may run one or more video-on-demand application(s) and a live-streaming application(s) which both make use of the client's in-network storage. The concurrently running applications may be running on different machines (e.g., multiple machines at the same home network) and may not directly communicate, except through user coordination

8.2. Per-Client Resource Policy

REQUIREMENT(S): A Provider MUST be able to specify resource policies (bandwidth share, storage quota, and network connection quota) to individual Consumers for reading from and writing data to its in-network storage within a particular range of time.

RATIONALE: Target Applications can rely on control of resources on a per-client basis. For example, application policy may indicate that certain remote clients have a higher bandwidth share for receiving data [LLSB08]. Additionally, bandwidth share for receiving data [LLSB08]. Additionally, certain data (e.g., chunks) may be distributed with a higher priority. As another example, when allowing a remote client to write data to a user's in-network storage, the remote client may be restricted to write less than 100MB of data in total. Since the client may need to manage multiple clients accessing its data, it should be able to indicate the time over which the granted resources are usable. For example, an expiration time for the access could be indicated to the DECADE-compatible server after which no resources are granted (e.g., indicate error as access denied).

8.3. Distributed Resource Sharing Specification

REQUIREMENT(S): A Provider MUST be able to indicate to its DECADE-compatible server, without itself contacting the server, resource control policies for a Consumer. The DECADE-compatible server MUST be able to authenticate the resource control policies.

RATIONALE: One important consideration for a DECADE-compatible server is scalability, since a single storage element may be used to support many users. Many Target Applications use small chunk sizes and frequent data exchanges. If such an application employed resource control and contacted the DECADE-compatible server for each data exchange, this could present a scalability challenge for the server as well as additional latency for clients.

The preferred way is to let requesting clients obtain signed resource control policies (in the form of a token) from the owning client, and then requesting clients can present the policy to a DECADE-compatible server directly. This can result in reduced messaging handled by the DECADE-compatible server.

8.4. Resource Set

REQUIREMENT(S): A DECADE-compatible server MUST allow specification on the following resources: storage, bandwidth, and number of connections, and MAY allow additional types of resources to be specified.

RATIONALE: The minimum set of resources need to include the most common resources.

9. Error and Failure Handling Requirements

9.1. Illegal Data Object

REQUIREMENT(S): A DECADE-compatible server SHOULD provide an error indicating that (1) data was rejected from being written, (2) deleted, or (3) marked unavailable.

RATIONALE: DECADE Storage Providers may require the ability to (1) avoid storing, (2) delete, or (3) quarantine certain data that has been identified as illegal (or otherwise prohibited). It is not specified how such data is identified, but applications employing DECADE-compatible servers should not break if a Storage Provider is obligated to enforce such policies. Appropriate error conditions should be indicated to applications.

EXCEPTION: A DECADE-compatible server should be able to be configured to suppress Illegal Data Object responses for security reasons.

9.2. Invalid Access Authorization

REQUIREMENT(S): A DECADE-compatible server SHOULD provide an error indicating that the request does not have a valid access authorization.

RATIONALE: DECADE-compatible clients may request data objects to which they do not have a valid authorization, and DECADE-compatible servers should be able to signal that this has occurred. Invalid authorization may be due to a combination of credential issues as well as additional policies defined by a Storage Provider.

EXCEPTION: A DECADE-compatible server should be able to be configured to suppress Invalid Authorization responses for security reasons.

9.3. Insufficient Resources

REQUIREMENT(S): A DECADE-compatible server SHOULD provide a response indicating to a DECADE-compatible client that resources (e.g., storage space) were not available to service its request (e.g., storage quota exceeded when attempting to write data).

RATIONALE: The Insufficient Resources response allows a client to back off, free up necessary resources or waiting for such resources to be freed.

EXCEPTION: A DECADE-compatible server may not provide such a response if doing so increases the load or due to security concerns.

9.4. Overload Condition

REQUIREMENT(S): A DECADE-compatible server, which is operating close to its capacity limit (e.g., too busy servicing other requests), MUST be permitted to reject requests and not be required to generate response to additional requests. A DECADE-compatible server MUST also be permitted to redirect requests as a load-shedding technique.

RATIONALE: The Insufficient Resources response allows a client to back off, free up necessary resources or waiting for such resources to be freed.

EXCEPTION: A DECADE-compatible server may not provide such a response if doing so increases the load or due to security concerns.

9.5. Attack Mitigation

REQUIREMENT(S): A DECADE-compatible server MUST be permitted to reject suspicious requests and not be required to generate responses (e.g., if a client's rate of requests exceeds a pre-defined threshold).

RATIONALE: Malicious clients may attempt to attack a DECADE-compatible server by specifying many chunks to increase total throughput or inciting overload conditions. A DECADE-compatible server is permitted to reject or ignore requests that are deemed suspicious according to policies set by its DECADE service provider.

10. Management Info Requirements

10.1. Account Status

REQUIREMENT(S): A Provider MUST be able to query the resource quota as well the current usage. The response from the server MUST include resource usage resulting from both the client's own usage and usage by other clients that have been authorized to read/write objects on that client's account.

RATIONALE: The resources used by a client are not necessarily directly-attached (e.g., disk, network interface, etc). Thus, the client cannot locally determine how much resources are being used. Before storing and retrieving data, a client should be able to determine which data is available (e.g., after an application restart).

11. Security Considerations

The security model is an important component of a DECADE-compatible system. It is crucial for users to be able to manage and limit distribution of content to only authorized parties, and the mechanism needs to work on the general Internet which spans multiple administrative and security domains. Previous sections have enumerated detailed requirements, but this section discusses the overall approach and other considerations that do not merit requirements.

11.1. Authentication and Authorization

A DECADE-compatible server must validate an request to access the in-network storage.

11.2. Confidentiality

DECADE-compatible Servers provide the ability to write raw data objects (subject to any policies instituted by the owner/ administrator of the Storage Provider). Thus, DECADE-compatible clients may opt to encrypt data before it is transported to the server.

11.3. Attack Mitigation

The particular resource control policy may be open to certain attacks by clients. For example by specifying many small chunks to increase total throughput or inciting overload conditions are techniques that may be used by a client.

12. IANA Considerations

There are no IANA considerations with this document.

13. References

13.1. Normative References

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

[RFC6646] Song, H., Zong, N., Yang, Y., and R. Alimi, "DECoupled Application Data Enroute (DECADE) Problem Statement", RFC 6646, July 2012.

13.2. Informative References

[I-D.ietf-decade-arch] Alimi, R., Rahman, A., Kutscher, D., and Y. Yang, "DECADE Architecture", draft-ietf-decade-arch-08 (work in progress), July 2012.

[LLSB08] Levin, D., LaCurts, K., Spring, N., and B. Bhattacharjee, "BitTorrent is an Auction: Analyzing and Improving BitTorrent's Incentives", SIGCOMM 2008, August 2008.

Appendix A. Acknowledgments

We would also like to thank Haibin Song for substantial contributions to earlier versions of this document. We would also like to thank Reinaldo Penno, Alexey Melnikov, Rich Woundy, Ning Zong, Roni Even, David McDysan, Borje Ohlman, Dirk Kutscher, Akbar Rahman, Xiao Zhu, Yunfei Zhang, Peng Zhang and Jin Peng for contributions and general feedback.

Authors' Addresses

Yingjie Gu
Huawei
No. 101 Software Avenue
Nanjing, Jiangsu Province 210012
P.R.China

Phone: +86-25-56624760
Email: guyingjie@huawei.com

David A. Bryan
Ethernot.org

Email: dbryan@ethernot.org

Yang Richard Yang
Yale University

Email: yry@cs.yale.edu

Peng Zhang
Tsinghua University/Yale University

Email: p.zhang@yale.edu

Richard Alimi
Google

Email: ralimi@google.com