

**CDNI Rate Pacing**  
**draft-caulfield-cdni-rate-pacing-02**

Abstract

Rate pacing is a class of network traffic shaping which limits the transmission rate of data over a network. This document defines CDNI extensions for downstream CDNs to support rate pacing on behalf of upstream CDNs.

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 [RFC 2119](#) [[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 30, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">CDNI Interfaces Impact</a>	<a href="#">3</a>
<a href="#">2.1.</a>	<a href="#">Footprint &amp; Capabilities Interface</a>	<a href="#">3</a>
<a href="#">2.2.</a>	<a href="#">Metadata Interface</a>	<a href="#">3</a>
<a href="#">2.2.1.</a>	<a href="#">RatePacing Metadata</a>	<a href="#">4</a>
<a href="#">2.3.</a>	<a href="#">Request Routing Redirection Interface</a>	<a href="#">4</a>
<a href="#">2.4.</a>	<a href="#">Logging Interface</a>	<a href="#">5</a>
<a href="#">2.5.</a>	<a href="#">Control Interface</a>	<a href="#">5</a>
<a href="#">3.</a>	<a href="#">Token Bucket Rate Pacing Algorithm and Parameters</a>	<a href="#">5</a>
<a href="#">3.1.</a>	<a href="#">TokenBucketParams Object</a>	<a href="#">6</a>
<a href="#">3.2.</a>	<a href="#">Token Bucket Metadata Example</a>	<a href="#">6</a>
<a href="#">4.</a>	<a href="#">IANA Considerations</a>	<a href="#">7</a>
<a href="#">4.1.</a>	<a href="#">CDNI Rate Pacing Algorithms Registry</a>	<a href="#">7</a>
<a href="#">5.</a>	<a href="#">Security Considerations</a>	<a href="#">7</a>
<a href="#">6.</a>	<a href="#">Acknowledgements</a>	<a href="#">8</a>
<a href="#">7.</a>	<a href="#">Normative References</a>	<a href="#">8</a>
	<a href="#">Author's Address</a>	<a href="#">8</a>

## [1.](#) Introduction

Rate pacing is a class of network traffic shaping which limits the transmission rate of data over a network. In the context of a Content Delivery Network (CDN), rate pacing provides an important business advantage to a Content Service Provider (CSP) by ensuring that a CDN which is delivering content on behalf of that CSP does not deliver significantly more data than necessary to an end client.

For example, suppose an end client is watching some Constant Bit Rate (CBR) video encoded at 1500 kbps. In the absence of rate pacing, the CDN delivering this content may send it to the client at 3000 kbps. If the client chooses to terminate the session before watching the entire video, up to half the transmitted data is wasted. This waste leads to unnecessary cost for the CSP and diminished useful capacity for the CDN.

Rate pacing requires configuration on a per-content basis. In order to enable rate pacing in a CDNI environment, the CDNI interfaces need to be extended to optionally support this feature.

This document describes:



1. CDNI interface extensions required for supporting rate pacing
2. a token bucket rate pacing algorithm for CDNs

## **2. CDNI Interfaces Impact**

### **2.1. Footprint & Capabilities Interface**

[I-D.ietf-cdni-footprint-capabilities-antics] defines the CDNI Footprint and Capabilities semantics. But at the time of writing, no FCI syntax specification has been accepted as a working group document.

[I-D.ietf-cdni-footprint-capabilities-antics] states that:

"The CDNI FCI specification SHOULD define the registry (and the rules for adding new entries to the registry) for the different capability types. Each capability type MAY further have a list of valid values. The individual CDNI interface specifications which define a given capability SHOULD define any necessary registries (and the rules for adding new entries to the registry) for the values advertised for a given capability type."

This document defines a new capability type: "RatePacing" to be added to the FCI capability types registry. The value of this capability contains one or more rate pacing algorithm names from the Rate Pacing algorithms registry ([Section 4.1](#)). For example, the value may be "token-bucket/v1" to indicate that the advertising CDN supports the token bucket algorithm described later in this document.

A CDN MAY advertise the "RatePacing" capability in the FCI if it implements this specification. A CDN MUST advertise the "token-bucket/v1" as a value in the list of algorithms if it advertises "RatePacing" as a capability.

### **2.2. Metadata Interface**

A new RatePacing metadata object is defined to represent the configuration for rate pacing. The RatePacing object has MIME type "application/cdni.RatePacing.v1". RatingPacing MAY appear within the metadata list of either HostMetadata or PathMetadata (i.e. may have either host-level scope or a path-level scope). The following section defines the properties of the RatingPacing object.



### **2.2.1. RatePacing Metadata**

The presence of the RatePacing Metadata indicates that a dCDN MUST comply with this specification in order to deliver a piece of content. The metadata indicates the rate pacing algorithm name required for delivering the content and the relevant parameters for that algorithm.

Property: algo

Description: Rate pacing algorithm name from the Rate Pacing Algorithms registry. Dictates the structure of the "params" value. For example, "token-bucket/v1".

Type: String

Mandatory-to-Specify: Yes.

Property: params

Description: An object containing algorithm-specific properties and values which are relevant to the rate pacing algorithm specified by the "algo" property. Each algorithm dictates its own parameters.

Type: Object

Mandatory-to-Specify: Yes.

### **2.3. Request Routing Redirection Interface**

The RRI is not impacted by rate pacing. However, if the metadata for a piece of content indicates that rate pacing is required by the uCDN, then a request router should only redirect requests for that content to CDNs which advertise "RatePacing" as a capability. The request router should also limit its choice of dCDNs to those which advertise the same rate pacing algorithm as is specified by the rate pacing metadata. Note that this behavior is not specific to rate pacing and is true of any CDNI feature.

For example, if the metadata for a piece of content includes a GenericMetadata object of type "application/cdni.RatePacing.v1" and the "algo" property in the value of that GenericMetadata is "token-bucket/v1", then the request router of the uCDN should only redirect requests for that piece of content to dCDNs which advertise a capability type of "RatePacing" and a capability value of "token-bucket/v1".



## **2.4. Logging Interface**

The rate at which a piece of content was delivered MAY be indicated via the LI. The "sc-rate" field indicates the rate in bytes per second as a decimal number. The bytes measured should correspond to the sc-entity-bytes field.

sc-rate:

format: DEC

field value: the average rate in bytes per second at which a response was delivered from Surrogate to client.

occurrence: there MUST be zero or exactly one instances of this field.

Note that existing fields defined by [[I-D.ietf-cdni-logging](#)] include the bytes delivered and the time taken to service a request, which could be used to estimate the delivery rate. However, the time taken includes the acquisition latency which is not relevant to rate pacing.

## **2.5. Control Interface**

The CI is not impacted by rate pacing.

## **3. Token Bucket Rate Pacing Algorithm and Parameters**

Token bucket is one example of a rate pacing algorithm. Token bucket is described generically by [[RFC1363](#)].

The token bucket algorithm is characterized by two parameters:

1. Rate - the number of tokens added to the bucket per second
2. Size - the maximum number of tokens in the bucket

This document specifies the CDNI Token Bucket Rate Pacing algorithm. It is based on the generic token bucket algorithm described above, but applied to a CDNI context.

For the purpose of this document, each token represents one byte transmitted as part of the body of an HTTP response from a Surrogate in a dCDN. Tokens do not represent bytes which are part of HTTP headers, the HTTP status line, TCP signaling, or any lower layer protocol.





The algorithm name "token-bucket/v1" is registered as a Rate Pacing algorithm. This algorithm name MUST appear as the value of the "RatePacing" capability. This name may also appear as the value of the "algo" property in the "RatePacing" metadata object.

If a RatePacing metadata object's "algo" value is "token-bucket/v1" then the metadata object's "params" MUST be an object of type TokenBucketParams, described below.

### **3.1. TokenBucketParams Object**

Property: rate

Description: Rate of tokens per second to be added to the bucket as described by the token bucket algorithm. This value MUST be a positive integer. Each token represents one byte.

Type: Integer

Mandatory-to-Specify: Yes.

Property: size

Description: Maximum number of tokens per bucket as described by the token bucket algorithm. This value MUST be a positive integer.

Type: Integer

Mandatory-to-Specify: Yes.

### **3.2. Token Bucket Metadata Example**

```
{
  "metadata": [
    {
      "generic-metadata-type": "application/cdni.RatePacing.v1",
      "generic-metadata-value": {
        "algo": "token-bucket/v1",
        "params": {
          "rate": 100000,
          "size": 25000
        }
      }
    }
  ]
}
```



#### 4. IANA Considerations

This document requests the following of IANA:

Addition of RatePacing in the CDNI Capability Registry defined in TBD.

Addition of "RatePacing" to the standard partition of the CDNI GenericMetadata Type Registry defined in [[I-D.ietf-cdni-metadata](#)]:

Type name	Specificati on	Version	MTE	STR
application/cdni.RatePacing	RFCthis	1	true	true
.v1				

Addition of "sc-rate" in the CDNI Logging Field Names Registry defined in [[I-D.ietf-cdni-logging](#)].

##### 4.1. CDNI Rate Pacing Algorithms Registry

IANA is requested to create a new registry, CDNI Rate Pacing Algorithms. The following table defines the initial values of the registry:

Algorithm Name	Specification
token-bucket/v1	RFCthis

New rate pacing algorithm registrations MUST specify RatePacing parameter objects as shown in [Section 3.1](#) and MUST describe the algorithm for rate pacing.

#### 5. Security Considerations

A malicious CSP might attempt to use rate pacing to instruct a dCDN to delivery some content at a very low rate thereby in order to exhaust the resources of a dCDN by forcing connection state to be maintained for longer than usual. The decision to enforce a rate is left to the discretion of a dCDN. An implementation of rate pacing should implement reasonable lower (and upper) bounds to avoid such cases.



## 6. Acknowledgements

The author would like to thank Francois Le Faucheur for his contributions and feedback.

## 7. Normative References

- [I-D.ietf-cdni-footprint-capabilities-semantic]  
Seedorf, J., Peterson, J., Previdi, S., Brandenburg, R.,  
and K. Ma, "CDNI Request Routing: Footprint and  
Capabilities Semantics", [draft-ietf-cdni-footprint-  
capabilities-semantic-03](#) (work in progress), July 2014.
- [I-D.ietf-cdni-logging]  
Faucheur, F., Bertrand, G., Oprea, I., and R.  
Peterkofsky, "CDNI Logging Interface", [draft-ietf-cdni-  
logging-14](#) (work in progress), October 2014.
- [I-D.ietf-cdni-metadata]  
Niven-Jenkins, B., Murray, R., Caulfield, M., Leung, K.,  
and K. Ma, "CDN Interconnection Metadata", [draft-ietf-  
cdni-metadata-07](#) (work in progress), July 2014.
- [RFC1363] Partridge, C., "A Proposed Flow Specification", [RFC 1363](#),  
September 1992.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate  
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

### Author's Address

Matt Caulfield  
Cisco Systems  
1414 Massachusetts Ave  
Boxborough, MA 01719  
USA

Phone: +1 978 936 9307  
Email: mcaulfie@cisco.com

