UltraBand CDR Guide

Version 5.7

Web: www.peerapp.com

PeerApp Inc.

29 Crafts Street
Newton, MA. 02458, USA

Support@peerapp.com

If you require any assistance, please contact us:
PeerApp Headquarters
PeerApp Inc.
29 Crafts Street
Newton, MA 02458, USA

http://www.peerapp.com
support@peerapp.com

[Document version: 5.7 - 1.00]

Table of Contents

Chapter 1: About This Guide

Overview

This document introduces the PeerApp Cache Detail Record (CDR) structure and parameters.

CDR's are files that provide detail on the traffic handled and delivered by the system. These files are used as an offline interface between the different PeerApp products.

CDR’s are CSV files with predefined structure. The CDR file name indicates its version number, where file structure and parameters can change between versions.

Chapter 2: CDR Structure

Cache Detail Record (CDR)

UltraBand generates CDRs that detail the content handled by the system. These records include information on cache in, cache out and forwarded content. In addition, each HTTP related CDR details information for its detected service group.

The following is a sample CDR record, and parameters description. For more examples, see CDR Examples, page 10.

10-03-10

22:34:04.397,HTTP,A834F7E5206B2E471765AAEE9C70F4148A323E17,

🔔www.example.com , /random.php?size=6000000&dev=500&rep=200 🔔,SESSION_ENDED,5997275,

10, 192.168.56.29:52333,192.168.57.27:80,22, 3, CACHE_OUT

Note: The ASCII bell character (hexadecimal value of x'07') differentiates a double quote, beginning or ending a string.

File Name

The CDR file name follows a predefined naming convention, indicating the CDR version. For example, pang_cdr.log.v43.ce-2.110913_234941.gz, as explained below.

Indicator	Description
pang_cdr.log	Determinative indicating it is a CDR file.
v43	Indicates the CDR version. In this example, version 4.3. **Note:** This determines the internal structure of the CDR file, which may vary between versions. See *"File Format" on page 2*.
ce-n	Cache engine from which the CDR file originated (n indicates the cache engine number).
032014_080003	Date and time, indicated by two digits (month, date, year _ hour, minutes, seconds)

File Format

CDR file format resembles that of the CSV.

1 All fields are ASCII characters.

2. Columns are separated by commas.
3. Nulls are represented by the character 0 (zero) placed between commas.
4. A bell character is always printed within the CDR. This is hard coded and cannot be configured to something else.
5. Line delimiter is new line (\n).

Chapter 3: CDR Parameters

Fields

The following table contains parameters needed for use in reading the CDR. For example: using the IPv6 sample shown below includes an embedded "(IDx)" after each field. This indicates the field ID from the table. These marks are not part of the CDR record and used here for reference only.

09-11-13 23:46:21.238(**ID1**),WEB_CACHE(**ID2**),crl.omniroot.com(**ID4**),/PublicSureServerSV.crl(**ID5**),RST(**ID6**),1448(**ID7**),0(**ID8**),41.188.8.175(**ID9**),52482(**ID10**),64.18.20.129(**ID11**),80(**ID12**),286(**ID13**),15(**ID14**),FORWARD(**ID15**),0(**ID16**),286(**ID17**),0(**ID18**),0(**ID19**),0(**ID20**),0(**ID21**),0(**ID22**),0(**ID23**),0(**ID24**),0(**ID25**),0(**ID26**),0(**ID27**),0(**ID28**),0(**ID29**),0(**ID30**),0(**ID31**),0(**ID32**),0(**ID33**),0(**ID34**),0(**ID35**),0(**ID36**),0(**ID37**),0(**ID38**),0(**ID39**),0(**ID40**),0(ID**41**),1(**ID42**),1448(**ID43**),51314(ID44),(ID45),595EBFF4(**ID46**),CDR_EOL(**ID47**)

ID	Name	Type	Definition	Length	Null	Comments
1.	Date Time	Date	This indicates the actual time frame.	22	No	Format DD-MM-YY HH24:MI:SS.MS
2.	Protocol	String	This is the format that is used for transmitting data from one device to another.	1-10	No	See "Protocols" on page 12.
3.	Hash	Hex string	The hash is used to generate a unique fixed length data output. This acts as a shortened reference to the original data.	40	Yes	E500FDF95C27ACE5C86068A6BE603A0BA857CBBB.
4.	Host	String	This is a physical network node that runs both the server and client programs. The host field value from the HTTP Request.	Up to 256	Yes	This only exists on HTTP and Netflix protocols. For example: xml.alexa.com

5.	Request URI	String	The HTTP Request - The URI value from the HTTP start line (web address).	Up to 1K	Yes	For example: /smartart/Text Effects/Textef fext None.swf
6.	Session Termination Reason	String	The "STR" details possible reasons for session termination.	Variable	No	See "Session Termination Reason" on page 13.
7.	Transferred Bytes	long	The transferred byte amount.		No	
8.	Session Duration	integer	The time measure of a session.		No	Seconds
9.	Source IP	4 octets separated by a dot(IPv4) 16 octets separated by a colon (Ipv6)1	The source IP is the subscriber IP address.	16	No	Pv4 or IPv6 addresses. 10.100.13.152 or 0001::2000:0000:0000:0002:01ce
10.	Source Port	Integer	The source port is the subscriber port.		No	
11.	Destination IP	4 octets separated by a dot (IPv4) 16 octets separated by a colon (Ipv6)	The destination IP is the endpoint IP address of the of the host.	16	No	IPv4 or IPv6 addresses.
12.	Destination Port	Integer	The destination port is the endpoint of the host.		No	
13.	HTTP Service Category ID	Integer	This is the first category index from the matched HTTP service. This is configured within the Policy Manager.		No	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.

14.	Http Service ID	Integer	This is the matched HTTP Service index. This is configured within the Policy Manager.		No	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
15.	CDR Type		A Cache/Content Data Record field which classifies by type. For more information, see "CDR Types" on page 13.			The field classifies the CACHE_OUT, VERIFY and VERIDFY_AT.
16.	Application Category ID	Integer	The matched Application Category index which is configured within the Policy Manager (Application). The Application field provides a high level summary of the content that can be used for reporting and policy creation, such as file sharing and software updates.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
17.	Service Category ID	Integer	The matched Service Category Index which is configured within the Policy Manager.		Yes	The Service Category ID is a legacy configuration based on the policy manager configuration.

18.	Internet Side Network Category ID	Integer	The matched Internet Side Network Category Index which is configured within the Policy Manager. This is associated to a sub-network address.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
19.	Subscriber Network Category ID	Integer	The matched Subscriber Network Category Index which is configured within the Policy Manager. The Subscriber Network is a category used to define the destination IP or subnet of the content.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
20.	Device Category ID	Integer	The matched Device Category Index which is configured within the Policy Manager. The User Agent is an HTTP header that specifies the type of browser/ device.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
21.	HTTP Referer Category ID (not used, always 0)	Integer	The matched HTTP Referer Category Index which is configured within the Policy Manager.		Yes	This field is not used. The default value of this field is 0.

22.	Content Type (MIME) Category ID	Integer	The matched Content Type (MIME) Category Index which is configured within the Policy Manager.		Yes	This field is not used. The default value of this field is 0.
23.	SmartFilter Category ID	Integer	The matched SmartFilter Category Index which is configured within the Policy Manager. This is used as a screen to filter/block specific content to the users.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
24.	Video Stream Bit Rate	Integer	This contains the video stream bit.		Yes	
25.	Video Stream Resolution	Integer	This contains the video stream resolution.		Yes	For YouTube traffic, the rate value is the itag URL parameter value that indicates the video resolution.
26.	Bandwidth Policy ID	Integer	This is a bandwidth policy index. It populates when a Rule with a bandwidth policy is matched.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
27.	Admission Policy ID	Integer	This is an Admission policy index. It populates when a Rule with an Admission policy is matched This is associated with blocking or re-directing sites.			The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.

28.	Cache Control Policy ID	Integer	This is a Cache Control policy index. It populates when a Rule with a Cache Control policy is matched. This policy disables/enables the media cache traffic to be cached in UltraBand.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
29.	Webcache Control Policy ID	Integer	This is a Web cache Control policy index. It populates when a Rule with a Web cache Control policy is matched. This policy disables the Web Cache traffic to be cached in UltraBand.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
30.	DSCP Policy ID	Integer	This is a DSCP policy index. It populates when a Rule with a DSCP policy is matched. This policy classifies cache out traffic in DSCP field that exists in the IPv4/IP6 headers.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
31.	Bandwidth Policy Rule ID	Integer	This is the matched rule index that applies the bandwidth policy on the current session.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.

32.	Admission Policy Rule ID	Integer	This is the matched rule index that applies the admission policy on the current session		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
33.	Cache Control Policy Rule ID	Integer	This is the matched rule index which applies the cache control (media cache control) policy on the current session.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
34.	Web-Cache Control Policy Rule ID	Integer	This is the matched rule index which applies the web-cache control policy on the current session.		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
35.	DSCP Policy Rule ID	Integer	This is the matched rule index which applies the DSCP policy on the current session		Yes	The values can be seen in UBView. This field is used by UBInsight for Traffic Analysis.
36.	Cache-In - (Number of Requests)	Integer	This field is used to specify the number of requests cached from in the traffic.		Yes	This is used in Web-Cache Sessions Only.
37.	Cache-In (Number of Bytes)	Integer	This field is used to specify the number of bytes cached from in the traffic.		Yes	This is used in Web-Cache Sessions Only.
38.	Cache-In (Duration)	Integer	This field is used to measure the amount of expired time of the incoming cache from the traffic.		Yes	This is used in Web-Cache Sessions Only.

39.	Cache-Out (Number of Requests)	Integer	This field is used to specify the number of cached -out requests from the traffic.		Yes	This is used in Web-Cache Sessions Only.
40.	Cache-Out (Number of Bytes)	Integer	This field is used to specify the number of cached -out bytes from the traffic.		Yes	This is used in Web-Cache Sessions Only.
41.	Cache-Out (Duration)	Integer	This field is used to measure the amount of expired time of the cache -out from the traffic.		Yes	This is used in Web-Cache Sessions Only.
42.	Forward (Number of Requests)	Integer	The process of forwarding requests directly to the requesting subscriber without caching it.		Yes	This is used in Web-Cache Sessions Only.
43.	Forward (Number of Bytes)	Integer	The process of forwarding the number of bytes directly to the requesting subscriber without caching it.		Yes	This is used in Web-Cache Sessions Only.
44.	Forward (Duration)	Integer	The time expended during the forwarding process.		Yes	This is used in Web-Cache Sessions Only.
45.	Referer	String	This is an http header field that identifies the address of the webpage linked to the resource being requested.		Yes	This is used in Web-Cache Sessions Only.

46.	Http Agent Hash		The http agent hash is used to generate the http agent's fixed-length output data that acts as a shortened reference to the original data			
47.	Full File Size	Integer	Displays available user requested file size in bytes. If the size is unknown, the CDR size is TDB			
48.	Cached File Size		Displays the amount of bytes cached from the requested user file.			
49.	Connection_ignore_reason		Displays the reason a connection state changed.			
50.	CDR_EOL	String	This is the end of the line for the CDR.		Yes	

Protocols

The following table details CDR protocols with examples:

SNMP ID	Name	Description
1	EDK	eDonkey
2	BT	Bittorent
3	GNUTL	Gnutella (not supported)
4	ARES	Ares
5	HTTP	HTTP
6	SILVERLIGHT	Silverlight
7	NETFLIX	Netflix
8	WEB_CACHE	Webcache

Session Termination Reason

The following table details possible reasons for session termination, as appears in CDR records.

ID	Name	Description
1	SESSION_ENDED	Session ended after entire file is delivered to the downloader.
2	CMDBMISS	Request range (in P2P) or file (HTTP) is not in the cache.
3	DOWNLOADER	Downloader closed the session
4	TIMEOUT	Session timed out.
5	ERROR	An error occurred serving the session.
6	RST	One of the sides reset the connection.
7	FIN	Connection terminated normally.
8	UPLOADER	Uploader closed the session.
9	UNKNOWN	None of the above reasons.
10	FORMAT_CHANGED	A change of the speed rate for Silverlight traffic.

CDR Types

The following table details the different CDR types.

ID	Name	Description
1	CACHE_OUT	Associated with cache-out.
2	CACHE_IN	A cache in transaction.
3	FORWARD	A forward transaction.
4	VERIFY	This is a verification process for a first time object access.
5	VERIFY_AT	This is a verification session that wasn't completed due to RST or TCP connection time-out from the client or server.

User-Agent Hash Mapping Files

The CDR records contain the User-Agent Hash. The mapping between the User -Agent Hash in the CDR file and the User-Agent String mapped through a User-Agent Database file. There is one User-Agent Database file is per Cache Engine.

> **Note:** The size of the User-Agent Database is limited a maximum 10000 entries, which is stored simultaneously. The User-Agent Database overwrites the least used agents with new entries once a day.

The following is the User-Agent database format:

[agent_hash], [agent_string]

The location of the files: /opt/pang/cdrs

The naming convention of the files are: http_agents.ce-[cache engine number].db

The following examples are not part of the CDR record and used here for reference only.

2898930B, Mozilla/5.0 (Windows NT 5.1; en-GB; rv:1.9.0.3) Gecko/2008092417 Firefox/3.0.3

CDA23AB3, Prey/0.5.3 (windows)

AB98F1B7, libhttp/1.61 (PlayStation 4)

8C5F7958, libhttp/3.01 (PS Vita)

2C968183,.Net 1.4.0.16 (Microsoft Windows NT 5.1.2600 Service Pack 3)

8BCC33ED,.Net 1.4.0.16 (Microsoft Windows NT 5.2.3790 Service Pack 2)

78614CA7,.Net 1.4.0.19 (Microsoft Windows NT 6.1.7601 Service Pack 1)

Chapter 4: CDR Examples

CDR Type Examples:

CACHE_OUT

09-11-13
23:44:54.355,WEB_-
CACHE,B5E0F73765E4F9D25BD50FD0C4604DAACD805000,cdn.niche.videosz.com,/48/300x250-
774633.jpg,RST,20311,0,197.149.12.69,3271,***.***.**.***,80,114,9,**CACHE_OUT**,109,114,0,0,111,0,0,0,0,0,0,0,11,13,0,26,26,26,26,26,0,0,0,1,19466,996,0,0,50044,http://feeds.videosz.com/custom/multi300x250/index.php?stid=10&ol=2&play=4&anim=1&sid=278&cid=593&a,3B73FE55,CDR_EOL

Note: For your reference, the CDR type is noted in **bold**.

CACHE_IN

09-11-13

23:46:22.134,HTTP,E4563542536A84761B1EE37B9A0FF2014EE5E500,video.thestat-icvube.com,/video/3/101c5117ab5-d374-43ac-8b82-df00f9aeb019.mp4,DOWNLOADER,1264692,23,41.188.38.216,64597,205.185.208.56,80,113,1,**CACHE_IN**,0,113,0,381,111,0,0,0,1171,0,0,0,0,0,0,0,0,0,0,0,,,,,,,,,,BD8D9C90,CDR_EOL

Note: For your reference, the CDR type is noted in **bold**.

FORWARD

09-11-13 23:46:21.408,HTTP,123C848BED162954F5CFBDC505CB4226C8A28D18,fr.download.nvidia.com,/Windows/331.65/331.65-desktop-win8-win7-winvista-64bit-international-whql.exe,SESSION_ENDED,3443854,130,197.149.49.37,49296,80.239.171.18,80,116,3,**FORWARD**,0,116,0,0,111,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,,,,,,,,,,E6C1BEA0,CDR_EOL

Note: For your reference, the CDR type is noted in **bold**.

VERIFY

09-11-13 23:46:21.677,HTTP,B50613095C3B6CB52E2CC5228613C6384B474CF3,www.helixstudios.net,/members/still_zips/hx103_scene45.zip,SESSION_ENDED,2654880,118,197.149.35.174,16647,216.201.82.10,80,213,11,**VERIFY**,0,213,0,382,111,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,,,,,,,,,,28B504D3,CDR_EOL

Note: For your reference, the CDR type is noted in **bold**.

VERIFY_AT

09-11-13 23:46:22.277,HTTP,0A2A400793E2F5B40420148C2E41478E0CB970AB,www.cycle-world.com,/wp-content/uploads/2012/12/Cooler-Boxer_010.jpg,DOWNLOADER,3953,31,197.149.5.23,55860,209.208.1.37,80,114,9,**VERIFY_AT**,109,114,0,0,111,0,0,0,0,0,0,0,11,13,0,26,26,26,26,26,,,,,,,,,,4466CB67,CDR_EOL

Note: For your reference, the CDR type is noted in **bold**.

CDR Protocol Examples

BitTorrent

14-11-13 09:00:49.678,**BT**,551C93A31984EF42EA495AA5BBD884840F7797E0,,,DOWNLOADER,12494514,19,83.223.186.231,49758,89.109.103.108,46765,0,0,CACHE_OUT,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,,,,,,,,,,00000000,CDR_EOL

Note: For your reference, the CDR protocol is noted in **bold**.

HTTP

09-11-13 23:49:41.241,**HTTP**,DDDC3E881DD1A31F1151DCD97DB17209E13F3A2D, 61.137.188.136:80 , /download/7F79C51D41CF2439C44ED2BF75268FE1BD06BA68/Ats-104.ts ,FIN,50740,30,197.149.19.225,52734,61.137.188.136,80,357,16,FORWARD,0,357,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,,,,,,,,,,33E8FE8F,CDR_EOL

Note: For your reference, the CDR protocol is noted in **bold**.

WEB_CACHE

09-11-13 23:44:54.355,**WEB_CACHE**,,t3.qlogo.cn,/mbloghead/6482cc2008129bb2f8fe/50,RST,0,0,197.149.6.162,26975,***.**.***.**,80,286,15,FORWARD,0,286,0,0,111,0,0,0,0,0,0,0,0,0,0,0,0,0,**0,0,0,0,0,0,1,0,51091,http://c.t.qq.com/i/19?pgv_ref=web.c.square.l2**,7F92EDDD,CDR_EOL

Note: For your reference, the CDR protocol is noted in **bold**.

Note: For your reference, please note the counters (36-45) that were added (also noted in **bold**).

YouTube Videoplayback Example:

14-11-13
12:29:48.368,HTTP,6A53245078C7C58D823FF3ADAF5D1D8700000000,r7---sn-nhpax-ua8l.googlevideo.com,/**videoplayback**?fexp=933508%2C937502%2C919319%2C914005%2C916611%2C901448%2C936912%2C936910%2C907231&id=0a819c53dad07035&burst=40&sver=3&ipbits=0&ip=77.127.74.78&gir=yes&signature=801E94AAE7BD46CCB83263FA6916EBA75D91635E.6EF2DE834FAC0A520C0D9788C4643D1DE2D5400D&mt=1384433429&mv=m&itag=135&ms=au&upn=22eThQPTKt0&sparams=algorithm%2Cburst%2Cclen%2Cdur%2Cfactor%2Cgir%2Cid%2Cip%2Cipbits%2Citag%2Clmt%2Csource%2Cupn%2Cexpire&expire=1384454308&clen=1448385&algorithm=throttle-factor&factor=1.25&source=youtube&dur=15.108&lmt=1384361440770122&key=yt5&ratebypass=yes,SESSION_ENDED,1187617,4,10.200.5.1,63451,212.199.205.210,80,54,7,CACHE_OUT,0,54,0,0,143,0,0,0,749,**135**,16,0,0,0,0,8,8,8,8,8,,,,,,,,,,2E0A163D,CDR_EOL

Note: For your reference, the CDR Youtube Videoplayback is noted in **bold**.

Note: The rate value is the itag URL parameter value which indicates the video resolution (noted in **bold**).