



Content Filtering Support Overview

This chapter provides an overview of the Content Filtering In-line Service feature.

This chapter covers the following topics:

- [Introduction, on page 1](#)
- [URL Blacklisting Support, on page 2](#)
- [Category-based Content Filtering Support, on page 6](#)
- [How URL Blacklisting and Category-based Content Filtering Work Concurrently, on page 17](#)
- [Content Filtering Server Group Support, on page 18](#)
- [External Storage, on page 19](#)
- [Bulk Statistics Support, on page 19](#)
- [Minimum System Requirements and Recommendations, on page 19](#)

Introduction

Content Filtering is an in-line service available for 3GPP and 3GPP2 networks to filter HTTP and WAP requests from mobile subscribers based on the URLs in the requests. This enables operators to filter and control the content that an individual subscriber can access, so that subscribers are inadvertently not exposed to universally unacceptable content and/or content inappropriate as per the subscribers' preferences.

The CF in-line service works in conjunction with the following products:

- GGSN
- HA
- IPSP
- PDSN
- P-GW

The Content Filtering service offers the following solutions:

- URL Blacklisting:

In the URL Blacklisting solution, all HTTP/WAP URLs in subscriber requests are matched against a database of “blacklisted” URLs. If there is a match, the flow is discarded, redirected, or terminated as configured. If there is no match, subscribers view the content as they would normally.

URL Blacklisting may/may not be a subscriber opt-in service, operators can enable URL Blacklisting either for all subscribers or for a subset of subscribers. Typical cases include applying a blacklisted database of child porn URLs to all subscribers so that they are inadvertently not exposed to such universally unacceptable content.

- Category-based Static Content Filtering:

In Category-based Static Content Filtering, all HTTP/WAP URLs in subscriber requests are matched against a static URL categorization database. Action is taken based on a URL's category, and the action configured for that category in the subscriber's content filtering policy. Possible actions include permitting, blocking, redirecting, and inserting content.

Typically Category-based Content Filtering is an opt-in service, subscribers self-choose a content-filtering policy or plan, such as Teen, Child, Adult, etc., and are subjected to content filtering as per their chosen plan. Also, the content filtering policies of different subscribers may be different, enabling differential access of content to them. This solution provides maximum flexibility, and is also referred to as the Policy-based Content Filtering.

Both URL Blacklisting and Category-based Content Filtering support can be concurrently enabled on a system.

Content Filtering uses Deep Packet Inspection (DPI) feature of Enhanced Charging Service (ECS) / Active Charging Service (ACS) to discern HTTP and WAP requests.

Qualified Platforms

CF is a StarOS in-line service application that runs on Cisco ASR 5500 platform. For additional platform information, refer to the appropriate *System Administration Guide* and/or contact your Cisco account representative.

Licenses Requirements

The URL Blacklisting, Category-based Content Filtering and External Content Filtering Server support through Internet Content Adaptation Protocol (ICAP) interface are licensed Cisco features. Separate feature licenses may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

For more information on ICAP feature, see the *ICAP Interface Support* appendix in the administration guide for the product that you are deploying.

URL Blacklisting Support

In the URL Blacklisting solution, a blacklist is a list of known URLs/URIs, which for some reason are being denied recognition. The blacklist can be obtained from a known source such as the National Center for Missing & Exploited Children (NCMEC, <http://www.missingkids.com>), Internet Watch Foundation (IWF) or any other IP source. The blacklist file is obtained from various sources in known formats and converted into a non human-readable optimized format (OPTBLDB) and then made available in the system. NCMEC database is provided in plain text whereas IWF database is provided in OPTBLDB format. For more information on the blacklist file, please contact your Cisco account service representative.



Important NCMEC database is provided by an external source, whereas IWF database is provided by Cisco.

Unlike the Category-based Content Filtering solution, which categorizes URLs as per a static database and takes different actions based on the different policies associated with subscribers, URL Blacklisting is applicable to all subscribers associated with a blacklisting-enabled rulebase. The same blacklist database is used for all subscribers, and for a specific URL, the same action is taken for all subscribers.

The blacklist file is downloaded and converted into a non human-readable optimized format (OPTBLDB) and then made available in the system. Once in place, all HTTP and WAP requests from subscribers are inspected in order to determine the requested destination URL/URI. If the URL/URI is not present in the blacklist then the request is passed on as usual. If the URL/URI is present in the blacklist, the request is dropped, or the flow is redirected or terminated as configured. There is no indication/messaging sent to the requesting subscribers that the requested HTTP/WAP URL/URI was rejected due to a blacklist match.

The blacklisting file can contain up to 32K URLs and the expected average size of blacklisting database (DB) is 2.5 to 5K.

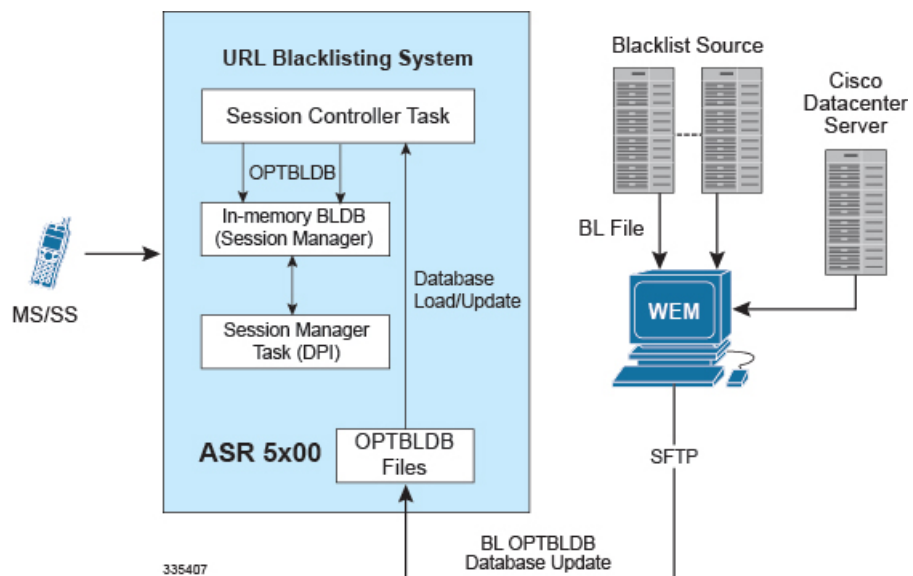
The URL Blacklisting match-method can be configured to either be generic or to look for any URL/URI in its exact, literal form.

The system generates usage/event data that can be utilized as the basis for blacklist reporting. The offline reports consist of, at a minimum, a running total of the number of times a match was made against the blacklist without any information regarding the specifics of the request.

The default/configured number of versions of the Blacklist database are maintained on the chassis (both the SPCs). This enables reverting to a particular version if required.

The following figure shows the high-level URL Blacklisting architecture with ECS, and other components in a deployment scenario.

Figure 1: High-Level Architecture of URL Blacklisting with ECS



URL Blacklisting Solution Components

The URL Blacklisting solution uses the deep-packet inspection capabilities of ECS for URL/URI extraction.

ECS functionality is managed by the following components:

- **Session Controller (SessCtrl):** The SessCtrl runs on the primary SPC/SMC and is responsible for managing ECS and URL Blacklisting services.
- **Session Manager (SessMgr):** A single SessMgr treats ECS charging and URL Blacklisting that is applicable to common subscriber sessions.

Apart from ECS, the URL Blacklisting solution uses the following components:

- Content Filtering Subsystem in ECS
- Web Element Manager (WEM)

Web Element Manager (WEM)

The WEM is a server-based application enabling complete element management of the system. The UNIX-based server application works with the network elements within the system using the Common Object Request Broker Architecture (CORBA) standard.

The WEM server must be set up with access to the following networks:

- Internet—to communicate with the source of the blacklist file (NCMEC/IWF/other)

The WEM application includes the following features:

- Single point of management for a large operator deployment
 - Service configuration and monitoring
 - Alarm/trap management for the WEM server
- URL Blacklisting database management functions:
 - Downloads the URL Blacklist database in OPTBLDB format from NCMEC or IWF, at configured schedule. The OPTBLDB file from NCMEC and IWF is merged and then pushed to the chassis.
 - Computes OPTBLDB suitable for updating the system.
- Distributes OPTBLDB files to the chassis automatically at configured interval.



Important

For more information on WEM, refer to the *Cisco Web Element Manager Installation and Administration Guide*.

How URL Blacklisting Works

This section describes how URL Blacklisting works.

Blacklist Updates

The following steps describe how the blacklist is updated in the system:

-
- Step 1** The WEM downloads the blacklist file from the specified source (NCMEC/IWF/other). NCMEC provides the file in clear text format which is then converted into a non-human readable optimized format (OPTBLDB), and IWF database provides the file in OPTBLDB format. The merged OPTBLDB file (NCMEC and IWF) is then pushed to the chassis.
- Step 2** The WEM pushes the optblk.bin file to the chassis (to /flash/bl) at pre-determined intervals. The optblk.bin file contains the full blacklist. If this file is verified to be correct, it replaces the optblk.bin file on the chassis and the last optblk.bin is rolled over.
- Step 3** The blacklist file is auto-detected by the Session Controller (SessCtrl), which verifies the integrity of the Blacklist database using checksums, and then loads it.
- The new blacklist is loaded only if it has been received properly. If the full Blacklist database is not found, corrupted, or if the loading fails, traps are generated. Correspondingly clear traps are also generated on a valid Blacklist database being available, and after a successful load.
- Step 4** The SessMgrs read the file and load the blacklisted URLs in a local in-memory database.
- Important** The URL Blacklisting feature is enabled only if the url-blacklisting action is set in any of the rulebases. Thus, the automatic detection of the Blacklist database, storing it in memory, and loading onto the SessMgrs will happen only if the url-blacklisting action is set in any of the rulebases.
- Step 5** The Blacklist database is loaded on each SessMgr as and when they come up (if URL Blacklisting is set in any rulebase) or when URL Blacklisting gets set in any of the rulebases.
- When the SessMgrs start for the first time or after recovery, if URL Blacklisting is set in any of the rulebases, the stored Blacklist database at SessCtrl is loaded onto the SessMgrs. This holds true for standby managers as well i.e., when standby managers come up, the Blacklist database is loaded onto them.
- Whenever a SessMgr is killed, standby manager which already has the Blacklist database loaded takes its place, and a new standby manager is created which loads the Blacklist database as part of SessMgr getting started for the first time.
- If SessCtrl is killed, while recovering it checks if URL Blacklisting is set in any of the rulebases, if set it will store the Blacklist database onto itself and load all the SessMgrs as well.
- Step 6** When a new Blacklist database is loaded on to the SessMgrs, the new database (and any stored versions that have rolled over) are synced to the other SPC so that after switchover, the proper Blacklist database can be accessed.
-

URL Blacklisting Action

The following steps describe how the URL Blacklisting feature works:

-
- Step 1** When an initial HTTP/WAP request comes for ECS processing and is processed by the ECS subsystem, a check is made to see if the URL Blacklisting support is enabled.
- Step 2** If enabled, the URL is extracted from the incoming request and is matched with the local in-memory Blacklist database. If a match is found for the URL in the Blacklist database, the packets are treated as per the blacklisting action configured—Discard, Redirect, or Terminate flow.
- In case of multiple HTTP requests in the same TCP packet, if any of the URLs match the packet is treated as per the blacklisting action configured.

If a match is not found, the request is allowed to pass through.

Category-based Content Filtering Support

The Category-based Content Filtering application is a fully integrated, subscriber-aware in-line service provisioned on chassis running HA services. This application is transparently integrated within the ECS, and utilizes a distributed software architecture that scales with the number of active HA sessions on the system.

Content Filtering policy enforcement is the process of deciding if a subscriber should be able to receive some content. Typical options are to allow, block, or replace/redirect the content based on the rating of the content and the policy defined for that content. For the list of content categories, refer to the *Category List* appendix in this guide.

Benefits of Category-based Content Filtering

The Category-based Content Filtering solution enables operators to ensure a simplified end-to-end traffic flow with a simple network topology. In-line deployment of Content Filtering provides a more attractive solution in contrast to out-of-line solutions where the filtering and policy enforcement is provided at some offload point that is decoupled from the bearer-processing layer.

The out-of-line model forces a session to make multiple hops through a redundant array of equipment which has a negative impact on traffic latency and limits subscriber and network visibility. In addition, the out-of-line model requires all subscriber sessions to be steered to the adjunct Content Filtering platform for policy enforcement regardless of whether this additional processing is needed. This leads to increased bandwidth provisioning requirements on gateway routers.

To facilitate network simplicity, it makes sense to leverage the benefits of deep packet inspection at a single policy enforcement point that is tied to the bearer processing layer. The advantages of this approach implemented include the following benefits:

- **Reduced processing latency:** In-line service processing eliminates unnecessary hand-offs and forwarding to external network elements.
- **Simplified policy provisioning:** Enables all policies like Content Filtering, ECS and QoS to be retrieved from same AAA/Policy Manager signaling interface thus reducing total volume of control transactions and associated delay.
- **Simplified provisioning and complete service integration:** Provisioning of separate resources like packet processing cards for processing subscriber data sessions and discrete services are eliminated. The same CPU can contain active Session Manager tasks for running Content Filtering and ECS charging.
- **Integration with Content Service Steering (CSS) architecture:** Enables applicable sessions to be forwarded to the in-line content filtering subsystem while delay and time sensitive voice/multimedia services immediately forwarded to Internet.
- **Service control:** Precise control over the interaction and service order handling of bearer flows with required applications like Content Filtering, ECS, Subscriber-aware Stateful Firewall, integrated Policy Charging and Rules Function (PCRF) for Service Based Bearer Control.

Apart from the advantages described previously, Category-based Content Filtering service reduces the requirement of over-provisioning of capacity at neighboring gateway routers. It also eliminates requirements of external Server Load Balancers and enhances the accuracy in subscriber charging records.

The Category-based Content Filtering solution has the following logical functions:

- Deep Packet Inspection (DPI) for Content Rating (event detection and content extraction)
- Content Rating Function with Static Rating of URLs
- Content Rating Policy Enforcement; for example, permit, discard, deny, redirect
- Content-ware accounting CF-EDR generation for events of interest



Important

Category-based Content Filtering can work in Static-only mode. Static-and-Dynamic and Dynamic-only Content Filtering modes are not supported.

ECS and Content Filtering Application

The Category-based Content Filtering subsystem is integrated within the Enhanced Charging Service (ECS) subsystem. Although it is not necessary to provision content-based charging in conjunction with content filtering, it is highly desirable as it enables a single point of deep-packet inspection for both services. It also enables a single policy decision and enforcement point for both services thereby streamlining the required number of signaling interactions with external AAA/Policy Manager servers. Utilizing both services also increases billing accuracy of charging records by insuring that mobile subscribers are only charged for visited sites content.

The Category-based Content Filtering solution uses Content Filtering Policy to analyze the content requested by subscribers. Content Filtering Policy provides a decision point for analyzed content on the basis of its category and priority.

The Category-based Content Filtering solution also utilizes ECS rulebases in order to determine the correct policy decision and enforcement action such as accept, block, redirect, or replace. Rulebase names are retrieved during initial authentication from the AAA/Policy Manager. Some possible examples of rulebase names include Consumer, Enterprise, Child, Teen, Adult, and Sport. Rulebase names are used by the ECS subsystem to instantiate the particular rule definition that applies for a particular session. Rulebase work in conjunction with a content filtering policy and only one content filtering policy can be associated with a rulebase.

For more information on rulebases and rule definitions, refer to the *Enhanced Charging Services Administration Guide*.

The CF Policy ID can be enabled depending on the subscriber (Child, Adult, etc.) and not based on the subscriber's device. This can be configured through PCRF for Gx using the **SN-CF-Policy-ID** attribute. For more information, refer to the *Subscriber Configuration* section in the *Content Filtering Service Configuration* chapter of this guide.



Important

If the **SN-CF-Policy-ID** sent by PCRF is 0 or junk (not configured on GGSN), then that value is ignored and the value of CF Policy ID remains to what it was before the PCRF message came (default, not set or any other value). For more information on this attribute, refer to the *AAA Interface Administration and Reference*.

The ECS subsystem includes L3–L7 deep packet inspection capabilities. It correlates all L3 packets with higher layer criteria such as URL detection within an HTTP header, it also provides stateful packet inspection for complex protocols like FTP, RTSP, and SIP that dynamically open ports for the data path.

The Content Filtering subsystem uses the deep-packet inspection capabilities of ECS for URL/URI extraction.

ECS functionality is managed by the following components:

- **Session Controller (SessCtrl):** The SessCtrl runs on the primary SPC/SMC and is responsible for managing ECS and Content Filtering services.
- **Session Manager (SessMgr):** A single SessMgr treats ECS charging and Content Filtering that is applicable to common subscriber sessions.

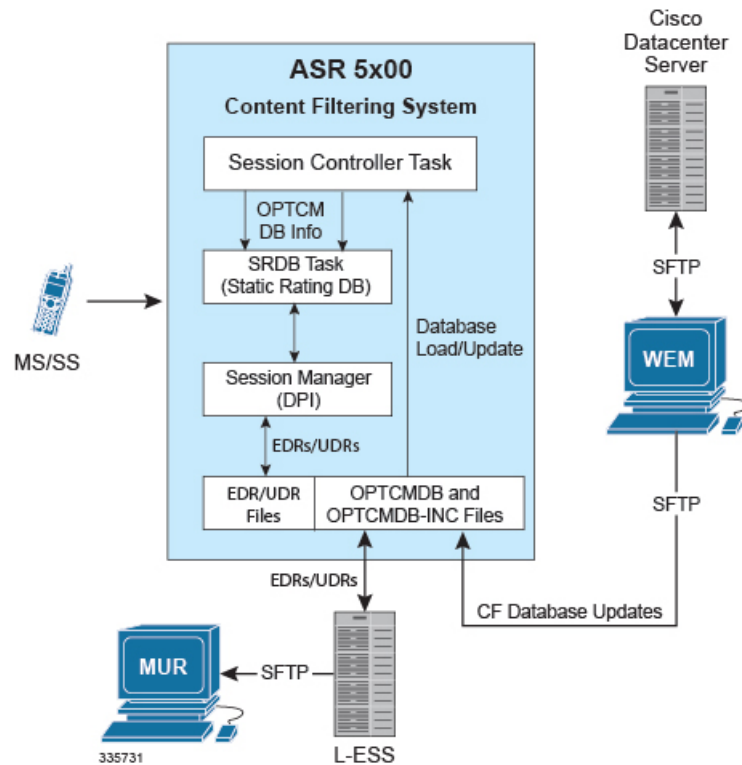
Components of Category-based Content Filtering Solution

The Category-based Content Filtering solution uses the following components:

- Content Filtering Subsystem in ECS
- Cisco Datacenter Server
- ECS Storage System (ESS)
- RADIUS Server/Policy Manager
- Content Rating Rules Update Server
- Web Element Manager (WEM)
- Mobility Unified Reporting (MUR) System

The following figure shows a high-level view of the Category-based Content Filtering architecture with ECS, and other components in a deployment scenario.

Figure 2: High-Level Architecture of Category-based Content Filtering



Category-based Content Filtering Subsystem

This is an internal categorization database (periodically synchronized with an external server) that provides ratings for publicly accessible traditional and mobile Web sites. When the SessMgr passes a URL/URI to internal list server, the list server returns a list of matching category ratings.

Static Rating Categorization Database (SRDB)

This is an internal categorization database (periodically synchronized with an external server) that provides ratings for publicly accessible traditional and mobile Web sites. When the SessMgr passes a URL/URI to internal list server, the list server returns a list of matching category ratings.

The list server is used to determine whether a Web site has already been classified. When the list server passes back a category rating to the filtering application, the rating is compared against the Category Policy ID applied for the subscriber to determine the appropriate action like accept, block, redirect, or replace. If the list server returns a clean rating, there is no need to perform a real-time analysis of any content delivered by the site.

The list server is used to determine whether a Web site has already been classified. When the list server passes back a category rating to the filtering application, the rating is compared against the Category Policy ID applied for the subscriber to determine the appropriate action like accept, block, redirect, or replace. If the list server returns a clean rating, there is no need to perform a real-time analysis of any content delivered by the site.

When a blocked or rejected content rating is returned, the SessMgr can insert data such as a redirect server address into the bearer data stream. If no rating is returned this means the site is capable of returning either clean or unacceptable content. In this case, the Content Filtering application uses the real-time dynamic analysis engine to examine additional content served by the site.

Each SRDB contains a replication object consisting of hash tables that map known Web sites and their subdirectories to their respective category ratings. The SessCtrl reads the index of SRDB tables with a data structure that associates keys with URL rating values and loads it onto the SRDB managers.

To boost performance and provide high availability, SRDB Manager provides functionality to load the Optimized Content Rating Master Database (OPTCMDB) volumes from its peer SRDB task. If the peer SRDB task is not in loading state then the OPTCMDB loading is done through SessCtrl to the recovered SRDB task.

DCCA Buffering Support

Static Content Filtering now interworks with DCCA buffering. Earlier, Static CF could buffer multiple packets at the same flow for rating and DCCA could handle buffering of single packet per flow. So, Static CF would not interwork with DCCA when DCCA buffering is enabled.

With the current implementation, CF does not send packets to DCCA after CF's rating, if DCCA has already buffered packets. When DCCA gets response of the buffered packet and has processed that packet, it will check if there are packets pending at CF to be processed, and will handle those packets one at a time. The remaining packets of the flow will be processed normally.

Content Rating Rules Update Server

This is a third-party content rating solution for exporting content filtering rules database information to the Category-based Content Filtering system. In addition, while exporting database updates, it collects reports of URLs processed by ECS and Content Filtering services that are reported as unknown in the deployed static rating database. This server analyzes these URLs and provides the rating in future updates for static rating database.

This server provides the following support to Cisco Datacenter Server for the content rating function:

- Provides full Vendor Format Master Database files (VFMDB) to Cisco Datacenter Server on request from Cisco Datacenter Server.
- Provides incremental Vendor Format Master Static URL Database file (VFMDB-INC) to Cisco Datacenter Server when any incremented VFMDB is available and requested from Cisco Datacenter Server.
- Receives the Unknown URLs file (Vendor Format Unknown Database File (VFUNKDB)) from Cisco Datacenter Server.

Cisco Datacenter Server

The Category-based Content Filtering solution provides a Cisco Datacenter Server to convert the VFMDB to SFMDB. It handles both full and incremental updates and processes them on a configured schedule.

This server is also responsible for distribution of SFMDB data files to WEM servers in the customer support infrastructure on a configured interval.

The server is responsible for following functionality as the Cisco Datacenter Server solution:

- Database fetching: Pulls VFMDB files from third-party Content Rating Server to Cisco Datacenter Server.
- Database conversion: Converts VFMDB files to SFMDB files. It also handles the incremented and unknown database files.
- Database poller: Provides the converted SFMDB database files for WEM in a preconfigured path.

- E-mail notification: Provides alerts and notification to the administrator for alarms.

External Storage

The external storage is a part of ECS Storage System in the ECS solution architecture.

The external storage is a storage application running on redundant highly available servers that collect and process EDRs and UDRs from which billing events and reports are generated. Either the system pushes the EDR/UDR files to the external storage, or the external storage fetches them from the system and processes them into formats suitable for billing mediation servers and MUR server. The external storage server consolidates the processed EDR/UDR files into a database for report generation through MUR. The database generated on an external storage by processing EDR/UDR records is a superset of the database required by MUR.

RADIUS Server and Policy Manager

The function of the RADIUS Server/Policy Manager in the Content Filtering solution is to provide per-subscriber Content Filtering provisioning information when a subscriber's session is established. It can also issue a Change-of- Authorization (CoA) to update an in-progress session to modify the Content Filtering policy for a subscriber.

The following are the basic functions provided by a RADIUS Server/Policy Manager in the Content Filtering solution:

- Support for the in/out ACL attributes to direct traffic through ECS for processing of subscriber traffic
- Support for ECS rulebase VSA to select the ECS rulebase to be applied to filtered traffic
- Support for Content Filtering Policy identifier VSA to select the content filtering policy within the selected rulebase for a subscriber
- Support exporting a subscriber provisioning record based on MSID to the customer service interface (Customer Care Interface) so that operator's customer care executive can see the provisioned content filtering policy for a subscriber

Web Element Manager (WEM)

The WEM is a server-based application providing complete element management of the system. The UNIX-based server application works with the network elements within the system using the Common Object Request Broker Architecture (CORBA) standard.

WEM server must be set up with access to the following networks:

- Internet: To communicate with the Cisco Datacenter Server which provides update files.

For Category-based Content Filtering, the WEM application includes the following features:

- Single point of management for a large Content Filtering Service operator deployment:
 - Content Filtering service configuration and monitoring
 - Alarm/trap management

- Configures and manages the operator-defined White/Black static rating database (WBLIST) for the network (WBLIST is maintained in SFMDB format).
For information on how to configure WBLIST, refer to the *Configuring WBLIST* section in this chapter.
- Content filtering database management functions:
 - Performs database processing in the background
 - Imports full and incremental SFMDB and SFMDB-INC files from the Cisco Datacenter Server on a configured schedule
 - Processes incremental SFMDB-INC updates from Cisco Datacenter Server maintaining an updated SFMDB file
 - Merge the operator's WBLIST database with the most recent SFMDB creating a SFCMDB
 - Computes an incremental update to the OPTCMDB-INC suitable for updating the Content Filtering subsystem that contains a previous version OPTCMDB
- Distributes OPTCMDB/OPTCMDB-INC files to the chassis automatically at configured interval. The initial file must be transferred manually for the automated push from WEM to work.

**Important**

For information on WEM, refer to the *Cisco Web Element Manager Installation and Administration Guide*.

Configuring WBLIST

Perform the following steps to configure WBLIST by the operator.

1. Create a WBLIST source file using the following command in a text editor:

```
vi wblast.txt
```

Enter the URLs to be categorized in the below format (URL category):

```
playboy.com DYNAM
```

```
colt.com DYNAM
```

```
dollarsgambling.com DYNAM
```

```
erowid.org DYNAM
```

```
sexetc.org DYNAM
```

2. Build **wblast.pub** file.

- WBLIST tool is available at `<EMS_Server>/server/tools/`.

- Generate the **wblast.pub** file as given below:

```
[root@pnqaextappsucs210-05 tools]# ./wblast -help
```

```
----- WHITE BLACK LIST Utility -----
```

```
-----
```

```
Usage :
```

```

-f <path> : path is wblast source file path
-o <output path> : WBLIST will be created at <path>/wblast.pub
: [default : current path]
-v <version> : WBLIST version number
: [default 1]
-h : Help

```

```

-----
[root@pnqaextappsucs210-05 tools]#

```

For example:

```

[root@pnqaextappsucs210-05 tools]# ./wblast -f wblast.txt -o
/export/home/SQA/ems/server/flash/cf/cfdatabases/wblastdb/ -v 2

```

```

----- WHITE BLACK LIST Utility -----

```

```

Setting WBLIST file path : /export/home/SQA/ems/server/flash/cf/cfdatabases/wblastdb//wblast.pub

```

```

-----
Source File : wblast.txt

```

```

WBLIST file path : /export/home/SQA/ems/server/flash/cf/cfdatabases/wblastdb//wblast.pub

```

```

-----
[root@pnqaextappsucs210-05 tools]#

```

3. This **wblast.pub** file automatically gets merged with the SFMDB if created at *<EMS Server>/server/flash/cf/cfdatabases/wblastdb/*.

Mobility Unified Reporting System

The Mobility Unified Reporting (MUR) application is a Web-based application providing a unified reporting interface for diverse data from the in-line service and storage applications. The MUR application provides comprehensive and consistent set of statistics and customized reports, statistical trending, report scheduling and distribution from chassis / in-line service product. For example, a subscriber's Quality of Experience, top 10 sites visited, top 10 users, and so on. The MUR application facilitates and enhances the operators' ability to simply and easily determine the health and usage of the network.

The MUR application supports the generation of various reports including CF-EDR reports in PDF and XML formats. The CF-EDR reports provide the summary of traffic over CF categories, CF actions, and CF ratings. It also provides the list of top N subscribers and URLs based on their unique subscriber's hit count and total usage.

- Summary Reports:
 - Category summary (volume/hits)
 - Action summary (volume/hits)
 - Rating summary (volume/hits)
- Top N Reports:

- Top N Subscribers by volume/hits
- Top N URLs by volume/hits

The CF-EDR files are pushed from external storage to MUR at a configured time interval and stored in a specified data directory on the MUR server. It can also create the files from CF-EDRs for unrated URLs which can be pulled by WEM.

For more information on the reports, refer to the *Mobility Unified Reporting System Online Help* documentation.

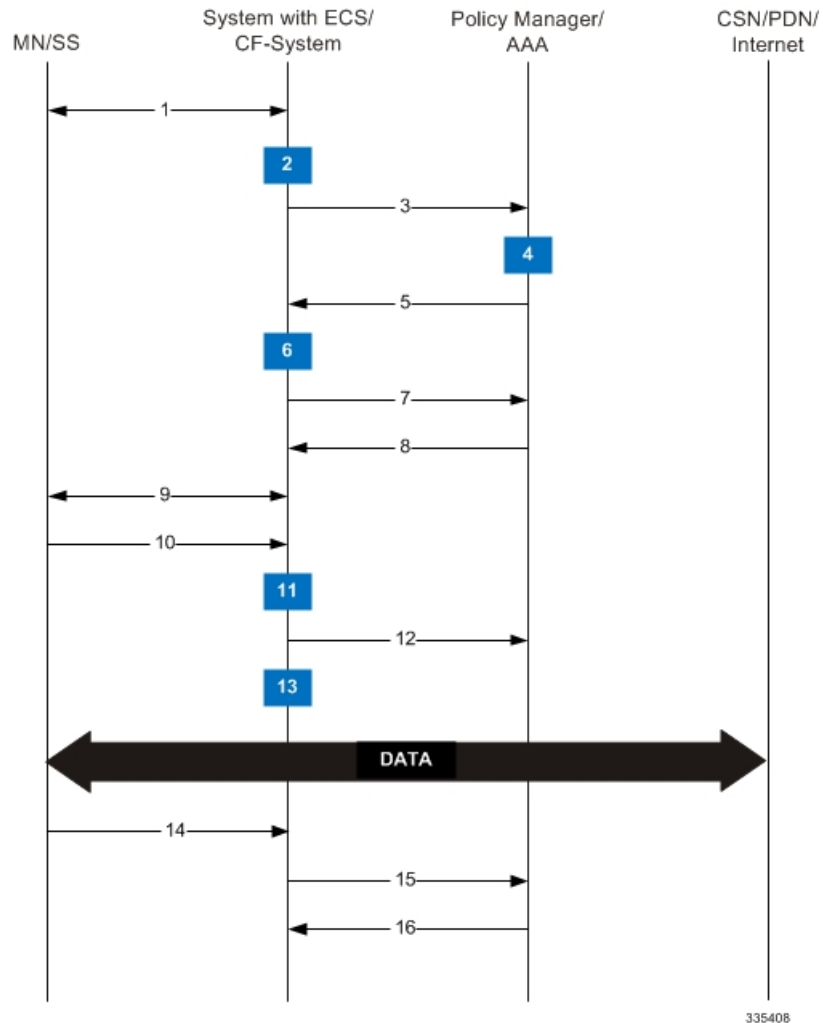
How Category-based Content Filtering Works

The Content Filtering Subsystem which is integrated into the ECS subsystem consists of an onboard static categorization database. The filtering service uses the Deep Packet Inspection (DPI) capabilities of the ECS subsystem to classify and partition application or protocol specific flows into virtual sessions.

Content analyzers are used to identify various types of flows such as HTTP, MMS/WAP, and POP3 E-mail. A typical HTTP request for a Web page, for example, invokes TCP and HTTP traffic analyzers. Any HTTP field including URLs or URIs can be identified. When a subscriber session is bound by CSS to an ECS running content filtering service, the URL/URI is extracted and compared against the static categorization database.

The following figure and the steps describe how Category-based Content Filtering works during a subscriber call:

Figure 3: Content Filtering Call Flow



- Step 1** MS requests for registration to the system.
- Step 2** System processes MS-related information with Content Filtering subsystem.
- Step 3** System sends the AAA Access Request to AAA server for MS.
- Step 4** AAA server processes the AAA Access Request from the Content Filtering subsystem to create the session, and the Policy Manager in AAA server uses subscriber identification parameters including NAI (username@domain), Calling Station ID (IMSI, MSID) and Framed IP Address (HoA) as the basis for subscriber lookup.
- Step 5** The Policy Manager and AAA generate and send an Access Accept message including all policy and other attributes to establish the session to the Content Filtering subsystem.

The Policy Manager and/or AAA include the following attributes in the Access Accept message:

- **Filter-Id or Access Control List Name:** Applied to subscriber session. It typically contains the name of the Content Service Steering (CSS) ACL. The CSS ACL establishes the particular service treatments such as Content Filtering, ECS, Traffic Performance Optimization, Stateful Firewall, VPN, etc. to apply to a subscriber session and the service order sequence to use in the inbound (SN-IP-In-ACL AVP) or outbound (SN-IP-Out-ACL AVP)

directions. Real-time or delay sensitive flows are directly transmitted to the Internet with no further processing required. In this case, no CSS ACL or Filter ID is included in the Access Response.

- **SN-CFPolicy-ID:** Applied to the subscriber content flow. Policy ID included in this attribute overrides the policy identifier applied to subscriber through rulebase or APN/Subscriber configuration. This content filtering policy determines the action to be taken on a content request from subscriber on the basis of its category. At anytime only one content filtering policy can be associated with a rulebase.
- **SN1-Rulebase:** This custom attribute contains information such as consumer, business name, child/adult/teen, etc.). The rulebase name identifies the particular rule definitions to apply. Rulebase definitions are used in ECS as the basis for deriving charging actions such as prepaid/postpaid volume/duration/destination billing and charging data files (EDRs/UDRs). Rulebase definitions are also used in content filtering to determine whether a type of user class such as teenagers should be permitted to receive requested content belonging to a particular type of category such as adult entertainment, gambling or hate sites. Rulebase definitions are generated in the Active Charging Configuration Mode and can be applied to individual subscribers, to domains or on per-context basis.

Step 6 Content Filtering subsystem creates a new session for MS.

Step 7 Content Filtering subsystem sends Accounting-Start messages to AAA server.

Step 8 AAA server sends Accounting-Start response message to Content Filtering subsystem.

Step 9 Content Filtering subsystem establishes data flow with MS.

Step 10 MS requests for data with URL name.

Step 11 Within the system access control list (ACL) processes the request and directs the request to ECS/Content Filtering subsystem based on the subscriber configuration.

Step 12 System performs ECS action on the content and then applies content filtering if required.

Within the system, if the bearer flow is treated by Content Filtering or other in-line services, the SessMgr feeds it to the Content Service Steering (CSS) API. If Content Filtering is the first service touch point, TCP and HTTP traffic analyzers within a given SessMgr utilize deep-packet inspection to extract the requested URL.

Step 13 The Content Filtering subsystem processes the URL access request.

When only Static Content Filtering is enabled, first the URL is looked-up in the cache maintained at SessMgr for static URL requests, if there is a hit, the category is returned, if its a miss, a URL look-up is performed by an onboard SRDB for static rating.

- If a category is returned, action is taken as configured for that category in the subscriber's Content Filtering policy:
 - allow: If the category is permitted by the subscriber's content filtering policy, the request is sent to the server, and the response transmitted to the subscriber's mobile.
 - content-insert: The system notifies the subscriber's mobile of the blocked content by inserting a specified message within the IP data stream, and prevents access to the requested content. The insert string is as specified in the subscriber's content filtering policy.
 - discard: The system silently discards the request packet(s).
 - redirect-url: The system inserts a specified redirect server address in the bearer data stream and returns an HTTP error message to the subscriber's mobile. The redirect address is as specified in the subscriber's content filtering policy.

If a new category (also referred to as x-category) is present in the OPTCMDB file, then action for all URLs for the new category can be configured at runtime. For information on how to configure x-category, refer to the *Configuring Content Filtering Policy* section in the *Content Filtering Service Configuration* chapter.

The redirect server may prompt the subscriber to send additional security credentials in order to access the requested content.

- **terminate-flow:** The system gracefully terminates the TCP connection between the subscriber and server, and sends a TCP FIN to the subscriber and a TCP RST to the server.
- **www-reply-code-and-terminate-flow:** The system terminates the flow with a specified reply code to the subscriber's mobile. The reply code is as specified in the subscriber's content filtering policy.
- If a category is not returned / the URL is not present in the database, the system takes the action as configured for the UNKNOW category in the subscriber's Content Filtering policy.
- If for the category returned there is no action configured in the subscriber's content filtering policy, the default action is taken.

Step 14 MS requests for session termination.

Step 15 System sends Accounting-Stop Request to the AAA server.

Step 16 AAA server stops the accounting for the MS for content filtering session and sends Accounting-Stop-Response to the system.

How URL Blacklisting and Category-based Content Filtering Work Concurrently

Both URL Blacklisting and Category-based Content Filtering can be concurrently enabled in a system. The following describes how URL blacklisting and content filtering are performed on HTTP/WAP traffic when concurrently enabled on a system:

Step 1 If both URL Blacklisting and Category-based Content Filtering are enabled, first URL blacklist matching is performed, and then, if required, content filtering is performed.

When an HTTP/WAP request comes for ECS processing, a check is made to see if the URL Blacklisting feature is enabled. If enabled, the URL is extracted from the incoming request and is matched with the local Blacklist database.

- If a match is found for the URL in the Blacklist database, the packets are subjected to the blacklisting action configured in the rulebase—Discard, Redirect, or Terminate flow. In case of multiple HTTP requests in the same TCP packet, if any of the URLs is blacklisted, then action is taken on the packet.
- If a match is not found in the Blacklist database, then Category-based Content Filtering is performed.

If Category-based Static Content Filtering is enabled, static rating is performed and action taken as configured for the category returned in the subscriber's content filtering policy.

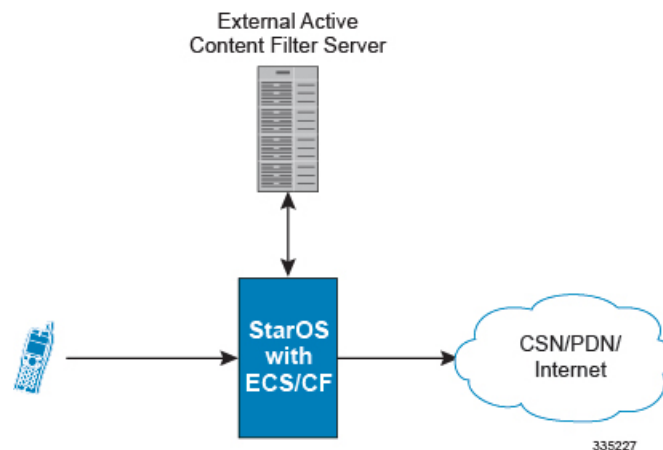
Step 2 If URL Blacklisting is enabled and Category-based Content Filtering is disabled, and a match is not found for the URL in the Blacklist database, the request is allowed to pass through, and no Content Filtering EDRs are generated for those flows.

Content Filtering Server Group Support

ECS supports the streamlined ICAP interface to leverage Deep Packet Inspection to enable external application servers to provide their services without performing DPI, and without being inserted in the data flow. For example, with an external Active Content Filtering (ACF) platform.

A high-level view of the streamlined ICAP interface support for external ACF is shown in the following figure.

Figure 4: High-Level View of Streamlined ICAP Interface with External ACF



The system with ECS is configured to support DPI and the system uses this capability for content charging as well. WAP and HTTP traffic is content filtered over the ICAP interface. RTSP traffic that contains adult content can also be content filtered on the ICAP interface. Only RTSP Request packets will be considered for content filtering over the ICAP interface.

If a subscriber initiates a WAP (WAP1.x or WAP2.0) or Web session, the subsequent GET/POST request is detected by the DPI function. The URL of the GET/POST request is extracted and passed, along with subscriber identification information and the subscriber request, in an ICAP message to the application server.

In the case of Category-based Content Filtering solution, the application server checks the URL on the basis of its category and other classifications like type, access level and content category and decides if the request should be authorized, blocked, or redirected by answering to the GET/POST with:

- A 200 OK message if the request is accepted.
- A 302 Redirect message in case of redirection. This redirect message includes the URL to which the subscriber should be redirected.
- A 403 Denied message if the request should be blocked.

Depending on the response received, the system with ECS will either pass the request unmodified, or discard the message, and respond to the subscriber with the appropriate redirection or block message.

Content Charging is performed by the ECS only after the request has been controlled by the application server. This guarantees the appropriate interworking between the external application and content-based billing. In particular, this guarantees that charging will be applied to the appropriate request in case of redirection, and that potential charging based redirections (i.e. Advice of Charge, Top Up page, etc.) will not interfere with the decisions taken by the application server.

The ACF performs the following functions:

- Retrieval of subscriber policies based on the subscriber identity passed in the ICAP message.
- Determining the appropriate action (permit, deny, redirect) to take for this type of content based on subscriber profile.
- Communication of the action (permit, deny, or redirect) decision for the URL back to the ECS subsystem.

For information on configuring the ICAP interface functionality for external ACF servers, see the *ICAP Interface Support* chapter of the *System Administration Guide*.

External Storage

External storage supports generation of EDR/UDR/FDR (xDR) files from the chassis. To store generated xDR files, on the Cisco chassis, the system allocates 512 MB of memory on the packet processing card's RAM. The generated xDRs are stored in CSV format in the */records* directory on the packet processing card RAM. These generated xDRs can be used for billing as well as for generation of reports to analyze network usage and subscriber trends. As this temporary storage space (size configurable) reaches its limit, the system deletes older xDRs to make room for new xDRs. Setting gzip file compression extends the storage capacity by approximately 10:1.

Because of the volatile nature of the memory, xDRs can be lost due to overwriting, deletion, or unforeseen events such as power or network failure or unplanned chassis switchover. To avoid losing charging and network analysis information, configure the CDR subsystem in conjunction with the external storage to offload the xDRs for storage and analysis.

For more information on the external storage, contact your Cisco Account representative.

Bulk Statistics Support

The system can be configured to collect bulk statistics (performance data) and send them to a collection server (called a receiver). Bulk statistics are statistics that are collected in a group. The individual statistics are grouped by schema. Content Filtering bulk statistics support only System schema.

The system supports the configuration of up to 4 sets (primary/secondary) of receivers. Each set can be configured with to collect specific sets of statistics from the various schemas. Statistics can be pulled manually from the chassis or sent at configured intervals. The bulk statistics are stored on the receiver(s) in files. The format of the bulk statistic data files can be configured by the user. Users can specify the format of the file name, file headers, and/or footers to include information such as the date, host name, chassis uptime, the IP address of the system generating the statistics (available for only for headers and footers), and/or the time that the file was generated.

For more information, see the *Configuring and Maintaining Bulk Statistics* chapter of the *System Administration Guide*.

Minimum System Requirements and Recommendations

This section identifies the minimum system requirements for components of the URL Blacklisting / Category-based Content Filtering solutions.



Important The hardware required for these components may vary, depending on the number of clients that require access, components managed, and other variables like EDR generation rate or CDR storage and processing requirements.

Certain basic server requirements are recommended for WEM and MUR system to exploit the CF solution. For information on these system requirements, refer to *Cisco Web Element Manager Installation and Administration Guide* and *Cisco Mobility Unified Reporting System Installation and Administration Guide*.

Cisco Datacenter Server System Requirements

This section provides information on the system requirements for Cisco Datacenter Server.



Important You must ensure that the minimum system requirements are met before proceeding with the Cisco Datacenter Server installation.

Hardware Requirements

The hardware required for the various components are given below:

- Dell PowerEdge 1950 server
 - 1.86 GHz Dual quad-core Intel Xeon CPU
 - 8 GB RAM
 - 2 * 146 GB RAID hard disk drive. The hard disk can be expanded up to 300 GB.
 - Gigabit Ethernet interfaces
 - CD-ROM Drive
- Operating Environment:
 - Red Hat Enterprise Linux 5.4

- or -

- Sun Microsystems Netra™ X4270 server
 - Quad-Core two socket Intel Xeon L5518 processor (1 * 4GB memory kit, 1333 MHz)
 - 32 GB RAM
 - 2 * 300GB 10K RPM SAS disks
 - SATA DVD drive
 - 8-port internal SAS HBA
 - Choice of AC or DC power supplies
- Operating Environment:

- Red Hat Enterprise Linux 5.4
- ZFS is the recommended file system with two ZFS pools.



Important

For the Cisco Datacenter Server 10.0 and earlier releases, it is recommended to use the hardware configurations of Dell PowerEdge 1950 server.



Important

For the Cisco Datacenter Server 11.0 and later versions, please use the hardware recommendations of X4270 server.

Additional Requirements on Chassis

The chassis requires the following additional hardware and memory to handle the Content Rating Master Databases; for example, for Category-based Content Filtering OPTCMDB. The memory required may vary with the size of rating databases used for content rating service.

- Minimum of two active packet processing cards are required
- Minimum 4 GB memory:
 - in Cisco chassis on Flash memory

