Intelligent Storage algorithms evaluation report

Deliverable 4.2

Avanzit Tecnología

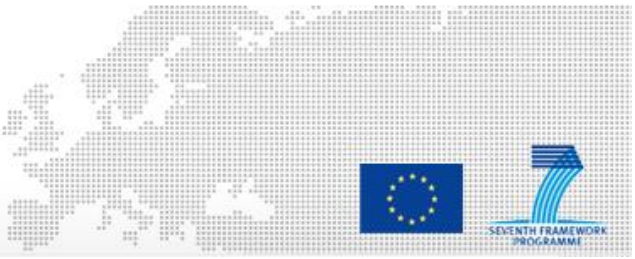

addpriv
Automatic Data relevancy Discrimination for a PRIVacy-sensitive video surveillance

Automatic Data relevancy Discrimination for a PRIVacy-sensitive video surveillance

SEC-2010.6.5-2 - Use of smart surveillance systems, data protection, integrity and sharing information within privacy rules

D4.2 – Intelligent Storage algorithms evaluation report

Due date of deliverable: 31.08.2012

Actual submission date: 11.09.2012

Start of project: 01 February 2011 Duration: 36 Months

Lead Contractor for this deliverable: Avanzit Tecnología

Revision: 1.0

Project co-funded by the European Commission within the Seventh Framework Programme		
Dissemination level		
PU	Public	
CO	Confidential, only for members of the consortium (including the Commission Services)	+

Revision History

Deliverable Administration and summary		
Project Acronym: ADDPRIV		Grant Agreement no: 261653
Document Identifier: ADDPRIV_20123108_WP4_AVANZIT_ISAER		
Leading partner: Avanzit Tecnología		
Report version: 1.0		
Report preparation date: 31-08-2012		
Classification: Confidential		
Nature: Other		
Author(s) and contributors: Avanzit Tecnología		
Status		Plan
		Draft
		Working
	+	Final
		Submitted
		Approved

The ADDPRIV consortium has addressed all comments received, making changes as necessary. Changes to the document are detailed in the change log table below.

Date	**Edited by**	**Status**	**Changes made**
27-08-2012	Avanzit Tecnología	First draft	
10-09-2012	Avanzit Tecnología	Working	Minor corrections
11-09-2012	Anova	Final	Quality check

Copyright

Citation

Avanzit Tecnología (2012). Deliverable 4.2 – Intelligent Storage algorithms evaluation report, www. ADDPRIV.eu

Acknowledgements

The work presented in this document has been conducted in the context of the EU Framework Programme project with Grant Agreement 261653 ADDPRIV (Automatic Data relevancy Discrimination for a PRIVacy-sensitive video surveillance). ADDPRIV is a 36 months project started on February 1st, 2011.

The project consortium is composed by: Anova IT Consulting (ANOVA), Kingston University Higher Education Corporation (KU), Politechnika Gdanska (GDANSK), Lancaster University (ULANCS), Avanzit Tecnologia, S.L. (AVANZIT), Hewlett Packard Italiana Srl (HP), Societá Per Azioni Esercizi Aeroportuali Sea SPA (SEA), Renfe Operadora (RENFE) and The Provost Fellows & Scholars Of The College Of The Holy And Undivided Trinity Of Queen Elizabeth Near Dublin (TCD).

More Information

Public ADDPRIV reports and other information pertaining to the project are available through ADDPRIV public website under www.ADDPRIV.eu

Table of contents

1. Introduction

1.1. ADDPRIV goals

The ADDPRIV project aims at improving public safety while ensuring the individuals' privacy rights. To achieve these goals, recorded images are classified as relevant or irrelevant according to the events detected by the Data Relevancy Discrimination algorithms [1] and the subsequent route reconstruction analysis [CITATION FOR RR NEEDED].

The relevancy or irrelevancy of the images is key in enhancing the privacy of the individuals since CCTV operators are only allowed access to the relevant video segments. Additionally, irrelevant footage can be safely erased [CITATION FOR SECURE ERASE NEEDED] which both further enhances the individuals' privacy and reduces the storage costs for the CCTV system owner.

1.2. Report aims

The Intelligent Storage algorithms are at the core of the system and will impact the usability of the whole system. This document describes the chosen algorithms and its suitability for the task.

1.3. Glossary

This section includes brief descriptions of abbreviations and uncommon terms used in this document.

- ***DRD*** – Data Relevancy Discrimination is an ADDPRIV system functionality, carried out by a set of algorithms which results in the video data categorization as relevant or irrelevant
- ***ED*** – Event Detection is an ADDPRIV system module which monitors video streams and automatically detects defined categories of threat related situations
- ***MP*** – Media Proxy is a subsystem which receives digital video from the existing CCTV system and provides synchronized frames to all the consumers modules
- ***PE*** – Privacy Enhancement is a subsystem which enables secure storage, deletion and access to archives of video data relevant to a given incident
- ***RR*** – Route Reconstruction is an ADDPRIV subsystem which analyses video frames collected by the cameras in the surveillance area and builds the route of a subject involved in a suspicious behaviour
- ***UI*** – User interface utilized by the end-user to interact with the developed system

2. System overview

This section presents a high-level architecture of the ADDPRIV system. In order to fully comprehend the system general function, the following definitions are required:

- Media Proxy (MP) is the module in charge of enriching and synchronizing the video frames proceeding from a multi-camera network with a timestamp used as unique ID for all the consumer ADDPRIV sub-modules: Event Detection (ED), Route Reconstruction (RR) and Privacy Enhancement (PE)
- Event Detection (ED) is the module in charge of processing video frames for detecting security relevant events.
- Route Reconstruction (RR) is the module responsible for processing in real time the input video frames by creating tagging and metadata used later by the Route Reconstruction algorithms.
- Privacy Enhancement (PE) is the module in charge of storing the video frames for later access from UI.

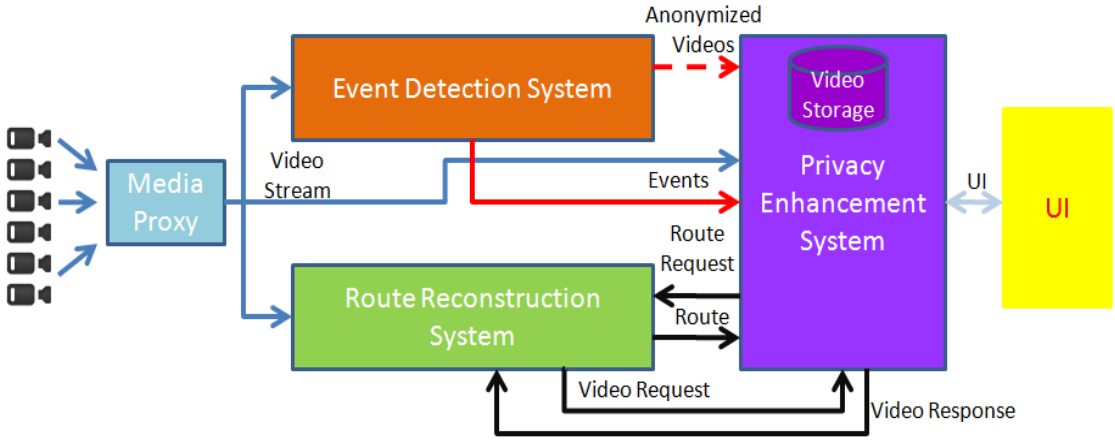

Fig. 1. The figure above presents the overview of the ADDPRIV architecture

The multi-camera network (in SEA or Renfe testing scenario) provides video frames to the MP system. When enriched by the MP, the video flux is consumed by ED and RR that constantly process video frames in real time and generate internal metadata and tagging information for finding and describing security relevant events (ED) and preparing metadata to be used for reconstructing route trees when requested by PE (RR).

When ED discovers a security relevant event, it sends an Event Descriptor to the PE; this information is permanently stored by the ADDPRIV system. The PE sends the Event Descriptor to the RR, which uses data from the object re-identification module that is designed to match the same object image representation across the various cameras that compose the multi-camera network.

The RR provides a Route Reconstruction tree related to the Route Request received from the PE that stores this information for its later provision only to authorized end-users via the UI.

The PE is in charge of securely deleting all the video frames that are older than 24 hours (or a configurable time interval) and that are not marked as security relevant event or part of a Route Reconstruction tree.

3. Intelligent Storage system overview

The Privacy Enhancement System (hereinafter PE) is the central component where the information generated and consumed by other subsystems is aggregated and leveraged.

The internal structures must be designed and developed to accommodate the needs of all these components.

In addition to orchestrating the interaction between all systems, the PE must do so while maintaining separate access level to the information and performing indexation, storage and retrieval of the received data efficiently.

For that purpose, the following subcomponents of the Privacy Enhancement System are foreseen:

- Data Management System: informs other components (such as the UI) and enforces the appropriate access level restrictions to the privacy-sensitive information.
- Events Database: persistent and distributed collection of events received from the ED. This database is specially constructed to allow efficient queries over the collected events.
- Routes Database: persistent and distributed collection of routes received from the RR. This database is specially constructed to allow efficient indexing of the routes and their association with events (in the Event Database) and video sequences (in the Video Database).
- Video Database: persistent store of the captured video stream distributed by each camera in the multicamera network, along with additional information including, at least, the relevancy of the frames from a security perspective and the capture date. It is designed to allow efficient access to individual frames and their associations with Events, Routes and Security Level according to heterogeneous usage patterns such as indexed access by camera ID and timestamp (in order to check/update relevancy of a given frame as the result of processing an Event or Route) or range scan by capture date and relevancy (in order to inform the Secure Deletion Agent of irrelevant and "old" frames or to construct a video stream to authorized users through the UI).

Together, these components constitute the **Intelligent Storage Subsystem,** which is aware of user activity and has access to sensitive information along all databases operating in a redundant, distributed cluster of commodity hardware.

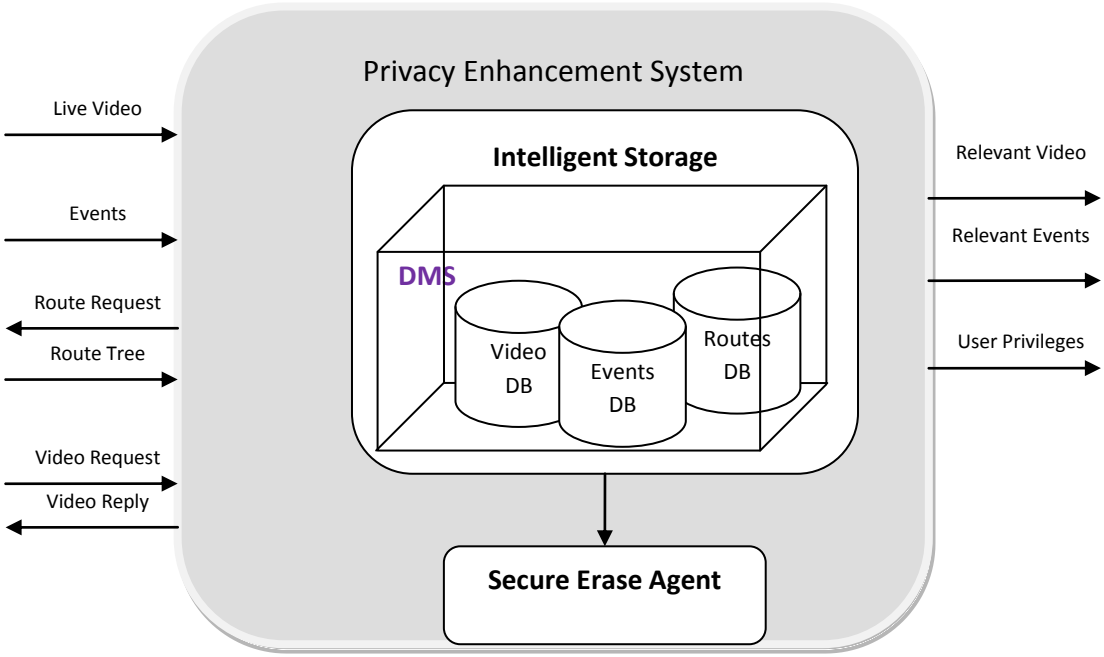

Figure 20: Privacy Enhancement System architecture

4. Intelligent Storage algorithms

As described in the previous section, Intelligent Storage module is comprised of a set of different components, which need to be orchestrated to work together in an efficient way to

1. Satisfy their intrinsic requirements, such as storage and indexing with the right levels of performance and correctness
2. Be suitable for consumers of the information, such as the User Interface

Each of these requirements will be described in turn and the relevant algorithms explained.

4.1. Indexes and DRD output

All indexes used to access the video data storage, as well as the Data Relevancy Discrimination output (events and reconstructed routes) is kept in a Mongodb[2] cluster, which ensures replication of data and high availability.

The document storage in Mongodb is especially well suited for the format of the messages defined in the ADDPRIV Solution Design [3].

In addition to document structure, Mongodb sports fast writes, measured up to 50 times faster than traditional databases, and supports ad-hoc queries, including the range queries needed to select all the frames in a time range.

4.1.1. Frame index

The frame index is a simple collection with the basic frame metadata: camera, timestamp of the capture, relevancy level and location of the frame image on the Video data storage.

The relevancy level is kept as a counter of the number of routes that cover the frame (include the frame in any of the non-discarded paths).

4.1.2. Video index

The video index is a simple collection storing the metadata for the generated videos: camera, time range of the capture, relevancy level, location of the video file and a reference to the route node that caused the video to be generated.

The reference to the route node is used to mark the video as irrelevant if the node itself is marked as such.

As for the frame index, the relevancy level is a counter of the number of routes that match the video. It will be greated than one in the unlikely event that two routes include a node with the same camera and time range.

4.1.3. Event database

Events triggered by the Event Detector are stored on this database as a collection.

The whole event, including extended (unknown) metadata not currently in use is stored.

The event database main role is serve as a starting point for route reconstruction request, but is also the main view on the system from the UI.

4.1.4. Routes database

Routes generated by the Route Reconstruction System are stored on this database as hierarchical documents, following the tree format described in the ADDPRIV Solution Design.

The routes in this database will be used for displaying to the CCTV operator in order to confirm its relevancy. This is required since RR uses statistical methods to build the routes.

4.2. Route pruning

Route pruning is the process by which a CCTV operator can review the results of the RR and confirm or reject the suggestion.

The algorithm for pruning is as follows:

- The CCTV operator is presented the most likely starting node for a route
- If the CCTV operator confirms it is relevant
    - All siblings of the node are pruned and associated videos' and frames' relevancy level will be decreased
    - The most likely next node is presented to the CCTV operator
- Else the CCTV operator rejects the suggestion
    - The suggested node and all its descendants are pruned and associated videos' and frames' relevancy level will be decreased
    - The next most likely sibling is presented to the CCTV operator
- Repeat until a single path remains

4.3. Video data storage

Footage received from the Media Proxy needs to be stored either permanently (if considered relevant) or for a given period of time (if irrelevant) until it is securely deleted.

The main goal of the ADDPRIV project being the enhanced privacy, video data storage is performed on a frame-by-frame basis. This comes at the cost of requiring additional effort to re-construct the video stream for presenting it to the User Interface, but allows fine-grained control over the produced relevant video.

Video data storage will be implemented on both cases on top of a standard filesystem provided by the underlying operating system. A Unix or unix-like system (such as BSD and its derivatives) allowing symbolic links is expected.

4.3.1. Frame storage

Frames are stored in a hierarchical filesystem structure, where the source camera for the stream is used as part of the hierarchy.

To achieve the desired real-time indexing speed for frames, the data can be stripped over several devices based on the source camera, allowing the Intelligent Storage to leverage the hardware buffers of each device when performing the frame writes.

Stripping by camera removes contention by stream reader processes trying to access the same device and enhances data locality, which will be useful when re-creating the video segments for relevant frames.

Frame files are named after their timestamp, so it is possible to partially recover the frame index in case of data migration or a catastrophic event.

All frames received are indexed by camera, timestamp and path to the actual image. They are initially marked irrelevant and thus made inaccessible by the DMS.

Stripping is achieved at the operating system level by using symbolic links to reference the camera subtree for each disk. The number of cameras to store per disk will need to be estimated to balance the size of video received from each camera to the write speeds of each drive plus a safe margin for reads and system operation. Attention should also be paid to disk size.

For the hypothetical case of a 1 gigabit/second network, fully saturated with video streams, and SSD drives allowing a moderate 200 megabytes/second writes, at least two drives would be needed.

4.3.2. Video generation

Videos are pre-generated whenever a relevant sequence is indexed, in preparation for them being available for the UI.

The Routes Database is used to determine the camera and time ranges that need to be generated. With that information, the frame index is used to identify all affected frames.

The Intelligent Storage system keeps a queue of video generation jobs and a pool of worker processes that perform them out-of-band. This allows the realtime processed to be prioritized and continue working without interruption.

Generated videos are added to the video index for fast lookup based on trigger event and reconstructed route node.

4.3.3. Video storage

Video storage algorithms follow the same structure as Frame storage, but it is independent in order to allow optimizing the drives for data reading (Frame storage is mostly used for writing, while the Video storage is mostly used for reads). Sharing the same drives is also possible.

4.4. Data Management System

The data management system controls access to the system and data, handling authentication and authorization of users.

4.4.1. User Authentication

Users authenticate against the system using their email address and password. Upon successful authentication, a unique API token is generated and can be used.

The authentication system is proactive against attacks by locking the accounts of users who have made too many invalid login attempts (actual number is configurable). The user can unlock the password by using his email address to receive the unlock url.

The user is also automatically logged out if he spends too many time idle (again, the actual value is configurable) to prevent piggyback attacks (an attacker using the open session).

4.4.2. User Authorization

User authorization is performed on all frame and video queries.

Three roles/access levels are defined:

- Level0: full access to all data. Intended for use by legally allowed parties, such as police carrying an investigation.
- Level1: only has access to data marked as relevant. Intended for CCTV operators.
- Level Irrelevant: only has access to data not marked relevant. Intended for the Secure Delete Agent.

A user is authorized to access the frame/video if her access level matches that of the content.

All queries are protected and no user has filesystem access to the actual data nor remote access to the databases. All access must be performed through an HTTP REST interface enforcing the DMS checks.

5. Functional testing

This section groups the test results of the automated test suite that accompanies the algorithms.

Tests are organized in features, which include a textual description. Each feature can have an optional background section, which defines the initial state of the system, and one or more scenarios, which describe specific test cases.

These background and scenario sections are built from one or more steps. Steps descriptions are parsed by the test engine and mapped to the actual algorithms' code.

5.1. Feature: The Intelligent Storage System will have at least three security levels for access to recorded video

The three levels will be corresponding to:

- Level-0: All the recorded information
- Level-Irrelevant: Privacy-sensitive video marked as irrelevant and thus eraseable
- Level-1: Privacy-sensitive video marked as relevant and made accessible to authorized personnel

Background

1. Given these recordings

camera_id	start_at	end_at	relevant
cam1	10:00am	10:15am	yes
cam2	10:00am	10:10am	no

2. And user '**police**' has role '**level0**'
3. And user '**eraser**' has role '**level-irrelevant**'
4. And user '**staff**' has role '**level1**'

Scenario: Level0 users can access relevant footage

1. Given user '**police**' requests footage from **cam1** from **10:00am** to **10:05am**
2. Then footage from **cam1** from **10:00am** to **10:05am** is provided

Scenario: Level0 users can access irrelevant footage

1. Given user '**police**' requests footage from **cam2** from **10:00am** to **10:05am**
2. Then footage from **cam2** from **10:00am** to **10:05am** is provided

Scenario: Level-Irrelevant users can not access relevant footage

1. Given user '**eraser**' requests footage from **cam1** from **10:00am** to **10:05am**
2. Then no footage is provided

Scenario: Level-Irrelevant users can access irrelevant footage

1. Given user '**eraser**' requests footage from **cam2** from **10:00am** to **10:05am**
2. Then footage from **cam2** from **10:00am** to **10:05am** is provided

Scenario: Level-1 users can access relevant footage

1. Given user '**staff**' requests footage from **cam1** from **10:00am** to **10:05am**
2. Then footage from **cam1** from **10:00am** to **10:05am** is provided

Scenario: Level-1 users can not access irrelevant footage

1. Given user '**staff**' requests footage from **cam2** from **10:00am** to **10:05am**
2. Then no footage is provided

5.2. Feature: Automatic annotation and 'relevancy indexing' of the images captured

The Intelligent Storage algorithms will store «the video files according to the classification they have received on the Data Relevancy Discrimitation Module»

Background

1. Given these cameras are recording

camera_id

cam1

cam2

2. And all the affected video has been indexed

Scenario: Video frames included in a reconstructed route are marked relevant

1. Given a route involving only **cam2** has been received
2. Then the frames included in all the possible routes should be marked as relevant

Scenario: Video frames not included in a reconstructed route are not modifed

1. Given a route involving only **cam2** has been received
2. Then no frame from **cam1** is marked relevant

Scenario: Video frames included in an event are marked relevant

1. Given an event on **cam1** has been received
2. Then the frames included in the event are marked relevant

Scenario: Video frames not included in an event are not modified

1. Given an event on **cam1** has been received
2. Then no frame from **cam2** is marked relevant

5.3. Feature: Read output from event detector

In order to be able to report the events to the UI and end users

As an Event Detector consumer

I want to receive, understand and index the detection events published

Scenario: Left luggage event notification

1. Given the received message

```
{  "type": "left luggage",  "cameraId": "main gate",  "timestamp": "2011-09-01 10:10:15", "bbox": { "top": 124, "left": 10, "bottom": 174, "right": 87 }, "custom_attribute": { "custom": "properties" }  }
```

2. When the UI asks for recent events
3. Then I should reply, at least, with the JSON message

```
{  "type": "left luggage",  "camera_id": "main gate",  "timestamp": "2011-09-01T10:10:15Z", "bbox": { "top": 124, "left": 10, "bottom": 174, "right": 87 }, "custom_attribute": { "custom": "properties" }  }
```

4. And include the internal id in the UI response

5.4. Feature: Cameras discovering

The service will connect to the Media Proxy (MP) and ask for the list of devices that are registered.

Scenario: Get video services

5. Given a URL

```
http://media_proxy/ver10/device/wsdl/GetDPAddresses
```

6. When make the request
7. Then will receive a list of active video services

```
{ {"camera_id": "main gate",  "TYPE": "IPv4" , "IP": "192.168.1.125"}, {"camera_id": "corridor 1",  "TYPE": "IPv4" , "IP": "192.168.1.137"} }
```

5.5. Feature: RTSP Sources

For each of the sources obtained, will begin a process that will store the information obtained.

Scenario: Video Storage

8. Given Video Source

```
{{"camera_id": "main gate",  "TYPE": "IPv4" , "IP": "192.168.1.125"}}
```

9. When make a RTSP request

rtsp://192.168.1.125/video

10. Then will recive a video streaming

```
Client – Server:        SETUP rtsp://192.168.1.125/video RTSP/1.0
                        CSeq: 2
                        Transport: RTP/AVP;unicast;client_port=4588-4589

Server – Client:    RTSP/1.0 200 OK
                        CSeq: 2
Session: 123124;timeout=60
Transport:RTP/AVP;unicast;client_port=4588-4589;
server_port=6256-6257
```

5.6. Feature: Frames Storage

Once the connection with RTSP video source has been established, each of the frames will be stored in a directory whose name is the camera id and will be named with the timestamp set by the media proxy.

Scenario: Video Storage

11. Given RTSP Connection

```
Client – Server:        SETUP rtsp://192.168.1.125/video RTSP/1.0
                        CSeq: 2
                        Transport: RTP/AVP;unicast;client_port=4588-4589

Server – Client:    RTSP/1.0 200 OK
                        CSeq: 2
Session: 123124;timeout=60
Transport:RTP/AVP;unicast;client_port=4588-4589;
server_port=6256-6257
```

12. When will obtain metainformation and frames
13. Then each frame will be storage in filesystem and will be named with the timestamp.

```
 sistemas@sol005:~/main gate$ ls -altr
total 8
drwxr-xr-x 12 sistemas sistemas 4096 2012-09-07 13:21 ..
-rw-r--r--  1 sistemas sistemas    1231 2012-09-07 13:23 1347024252.jpg
-rw-r--r--  1 sistemas sistemas    1315 2012-09-07 13:23 1347024253.jpg
-rw-r--r--  1 sistemas sistemas    1230 2012-09-07 13:23 1347024254.jpg
-rw-r--r--  1 sistemas sistemas    1201 2012-09-07 13:23 1347024255.jpg
-rw-r--r--  1 sistemas sistemas    1223 2012-09-07 13:23 1347024256.jpg
-rw-r--r--  1 sistemas sistemas    1321 2012-09-07 13:23 1347024257.jpg
-rw-r--r--  1 sistemas sistemas    1005 2012-09-07 13:24 1347024258.jpg
-rw-r--r--  1 sistemas sistemas    1358 2012-09-07 13:24 1347024259.jpg
-rw-r--r--  1 sistemas sistemas    1231 2012-09-07 13:24 1347024260.jpg
-rw-r--r--  1 sistemas sistemas    1231 2012-09-07 13:24 1347024261.jpg
-rw-r--r--  1 sistemas sistemas    1231 2012-09-07 13:24 1347024262.jpg
-rw-r--r--  1 sistemas sistemas    1248 2012-09-07 13:24 1347024263.jpg
```

```
drwxr-xr-x  2 sistemas sistemas 4096 2012-09-07 13:24 .
```

5.7. Feature: Create Reconstruction Route Videos

Upon receiving the route request due to suspicious activity, the system composed the videos of each of the cameras that are involved in a given time period.

Scenario: Publish RSTP

14. Given Petición de ruta

```
{ "requestId": "rq-0912",
  "cameraId": "cam-12",
  "timestamp": "Mon, 01 Aug 2011 10:57:59 GMT",
  "bbox": { "top": 120, "left": 50, "bottom": 175, "right": 97 }
  “RRRoot”:
   {
"requestId": "rq0912",
  "routes": [ {
   "pastTree": {
    "subjectId": "rr-sub-123",
    "cameraId": "cam2",
    "bboxes": [
     { "timerange": ["Mon, 01 Aug 2011 10:57:59 GMT",
                "Mon, 01 Aug 2011 10:58:01 GMT"],
      "top": 120, "left": 50, "bottom": 175, "right": 97
     },
     { "timerange": ["Mon, 01 Aug 2011 10:58:01 GMT",
                "Mon, 01 Aug 2011 10:58:07 GMT"],
      "top": 170, "left": 70, "bottom": 195, "right": 117
     }
    ],
}
}
```

15. When We generated video frames involved in that period.
16. Then will get the video and your name.

```
VlcXX7WAs8vm5IMWBRoJEA

sistemas@sol005:~/main gate/videos$ ls -altr
total 8
drwxr-xr-x 12 sistemas sistemas 4096 2012-09-07 13:21 ..
-rw-r--r--  1 sistemas sistemas    14587 2012-09-07 13:42 VlcXX7WAs8vm5IMWBRoJEA.mp4
drwxr-xr-x  2 sistemas sistemas 4096 2012-09-07 13:24 .
```

17. And Publishing Video

```
rtsp://sol005:5556/ VlcXX7WAs8vm5IMWBRoJEA.mp4
```

6. References

[1] GDANSK (2012) Deliverable 3.6 – Data Relevancy Discrimination algorithms evaluation report

[2] Mongodb is an open source document-oriented database system. Instead of storing data in tables as is done in a "classical" relational database, MongoDB stores structured data as JSON-like documents with dynamic schemas – www.mongodb.org

[3] HP (2011) Deliverable 2.3 – ADDPRIV solution design