

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

Vigilant Guard: Giving AI Eyes to ZoneMinder Without Handing Your Home to the Cloud

Publicado em 2026-05-29 12:54:27



FACT BOX

- **Solution name:** Vigilant Guard.
- **Purpose:** detect humans in ZoneMinder events using local AI.
- **Main engine:** YOLOv8 running on CPU.

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

compatible clients.

- **Facial recognition:** optional, using a MySQL/MariaDB database.
- **Philosophy:** no mandatory cloud, no external dependency, full local control.

Vigilant Guard: Giving AI Eyes to ZoneMinder Without Handing Your Home to the Cloud

Some technological solutions are born from grand strategic plans. Others arise from a very simple need: a camera detects movement, but the system must know whether it is seeing a person, a shadow, a cat, a leaf in the wind, or merely another false positive with dramatic ambitions.

Vigilant Guard was born from that practical need: to add to **ZoneMinder** a local artificial intelligence layer capable of analysing motion events, detecting humans and sending

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

to do what it does well — capturing video, detecting motion and creating events. Vigilant Guard then steps in as a second observer, more selective, more attentive and less easily impressed by nocturnal shadows.

1. The classic problem of domestic video surveillance

Anyone who uses video surveillance systems knows the dilemma: if sensitivity is too low, the system misses important events; if sensitivity is too high, it starts firing alerts for everything and nothing. A cloud passes in front of a light, a tree moves, an insect crosses the lens, and the system enters apocalypse mode.

ZoneMinder is a robust and flexible tool, but traditional motion detection is mainly based on image changes. That is enough to understand that “something changed”, but not enough to understand **what changed**.

That is exactly where Vigilant Guard enters the picture: it does not replace ZoneMinder; it adds discernment.

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

Everything runs inside the user's own infrastructure.

```
ZoneMinder
```

```
↓
```

```
Event filter
```

```
↓
```

```
notify-image.sh
```

```
↓
```

```
UNIX socket /tmp/humans_socket
```

```
↓
```

```
Vigilant Guard
```

```
↓
```

```
YOLOv8
```

```
↓
```

```
optional face_recognition
```

```
↓
```

```
optional MySQL/MariaDB
```

```
↓
```

```
ntfy / crops / alerts
```

ZoneMinder creates the event. A filter executes the `notify-image.sh` script. That script sends the event path to a UNIX

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

The chain is short, clear and debuggable. And, like all good UNIX systems, when it fails, it leaves clues — provided we have enough logs and the patience of a craftsman.

3. The role of the UNIX socket

One of the most important design decisions was to use a local **UNIX socket** instead of HTTP, REST APIs or another heavier mechanism. The socket is fast, simple and well suited for communication between processes on the same machine.

The socket path used is:

```
/tmp/humans_socket
```

There was, however, a critical detail: ZoneMinder runs the filter as the `www-data` user, while the Python service runs as `root`. If the socket is created as `root:root` with restrictive permissions, the ZoneMinder script cannot deliver events.

The solution was to create the socket with the `www-data` group:

```
srw-rw---- 1 root www-data ... /tmp/humans_s
```

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

4. Human detection with YOLOv8

The main detection engine is **YOLOv8**, using the lightweight `yolov8n.pt` model. This choice was deliberate: the system runs on a virtual machine without a dedicated GPU, so the solution had to work on CPU.

A typical configuration is:

```
HUMANS_YOLO_MODEL=/opt/VigilantGuard/models/  
HUMANS_MAX_EVENT_FRAMES=4  
HUMANS_MAX_IMAGE_WIDTH=960  
HUMANS_PERSON_CONF=0.35  
HUMANS_MIN_PERSON_AREA_RATIO=0.003
```

The `HUMANS_PERSON_CONF` parameter controls the minimum confidence required from the model. Lower values increase sensitivity, but may also increase false positives. Higher values reduce noise, but may miss people in dark, partial or distant images.

The goal is to find a balance: detect well enough without turning every shadow into a domestic emergency.

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

person is clearly visible. For that reason, Vigilant Guard gives priority to two essential images:

- `alarm.jpg` — usually the image associated with the alarm moment.
- `snapshot.jpg` — the event reference image.

The logic was adjusted to ensure that both images are analysed whenever they exist. Only afterwards are additional event images selected across the sequence.

This improves the probability of detecting a person, especially when the person passes quickly in front of the camera.

6. Optional facial recognition

After detecting a person, the system can attempt to identify the face using the `face_recognition` library, comparing facial encodings with a MySQL or MariaDB database.

The table may have a simple structure:

```
CREATE TABLE faces (  
    id INT AUTO_INCREMENT PRIMARY KEY,
```

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

);


Facial recognition is, however, optional. The solution remains useful without it: it detects humans, crops the figure and sends an alert. Named identification only happens when there is a visible face, with sufficient quality, previously registered in the database.

In practice, YOLO answers the question: “**Is there a person?**”. Facial recognition then tries to answer the next question: “**Who is it?**”.

7. Alerts via ntfy

Alerts are sent through **ntfy**, a simple and lightweight solution particularly well suited for notifications in self-hosted systems.

A typical message may indicate:

```
 Human figure detected – http://server/up1
```

Or, when a face is recognised:

```
 Face identified: Name – http://server/up1
```



8. systemd service: turning a script into a sentry

To run continuously, Vigilant Guard operates as a systemd service:

```
sudo systemctl status Vigilant-Guard.service
sudo journalctl -u Vigilant-Guard.service -f
```

The service creates the socket, loads known faces, initialises the YOLO model and waits for event paths sent by the ZoneMinder script.

Using `systemd` allows automatic restart, CPU limits, lower scheduling priority and clean integration into the system boot process.

9. Limits and tuning

Vigilant Guard does not solve every video surveillance problem by itself. If ZoneMinder does not create an event, Vigilant Guard will not be called. The first link in the chain remains the configuration of zones, sensitivity, minimum area and motion criteria inside ZoneMinder itself.

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

- **In Vigilant Guard:** YOLO confidence, minimum person area, number of images analysed per event, maximum image width.

If the system detects too many people — for instance, pedestrians in the street — the first solution should be to adjust the zones in ZoneMinder. It is better to prevent an unwanted event from being created than to ask AI to compensate for poor surveillance geometry.

10. A local, simple and publishable project

Vigilant Guard was born from a concrete need, but it has the potential to be useful to other ZoneMinder users. Many people want to add human detection to their cameras, but do not want to depend on closed platforms, subscriptions, mandatory cloud services or commercial black boxes.

This solution follows a different philosophy: local software, explicit configuration, readable logs, simple integration and full control by the user.

It does not aim to be a universal intelligent surveillance platform. It aims to be a clear, small and effective tool: receive events, analyse images, detect humans and send alerts.

Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.

artificial intelligence to ZoneMinder without dismantling the existing architecture. Instead of replacing the video surveillance system, it adds a layer of interpretation.

The result is a more attentive digital watchman: it does not merely see movement, but tries to understand whether that movement corresponds to a person.

And perhaps that is the best definition of this small solution: a **wide-awake night watchman**, powered by Python, YOLO, UNIX sockets and a certain amount of technical stubbornness. The good kind of stubbornness — the artisanal kind that turns small problems into useful systems.

Author: Francisco Gonçalves

Technical and editorial co-authorship: **Augustus Veritas.**

Technical article published within the experimental technology, open software, local artificial intelligence and independent infrastructure projects of Fragmentos do Caos.


Blogue Fragmentos do Caos



A verdade nasce onde o pensamento é livre.



Fragmentos do Caos: [Blogue](#) • [Ebooks](#) • [Carrossel](#)

 Esta página foi visitada ... vezes.

[Contactos](#)