

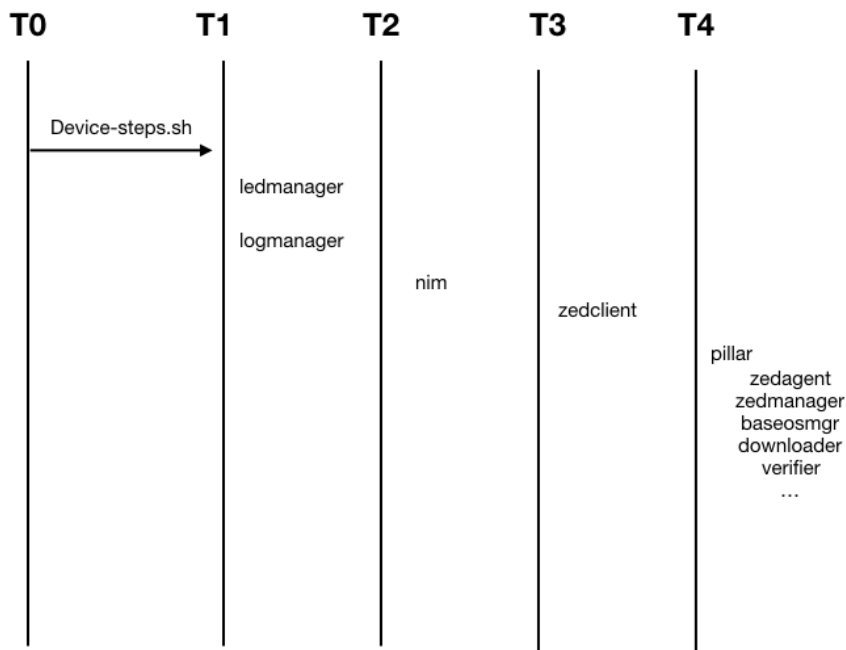
NodeAgent: EVE node health monitor module

Problem Statment

Currently, on EVE software module, ZedAgent is responsible for top level orchestration, basos upgrade validation, cloud connectivity for configuration /status.

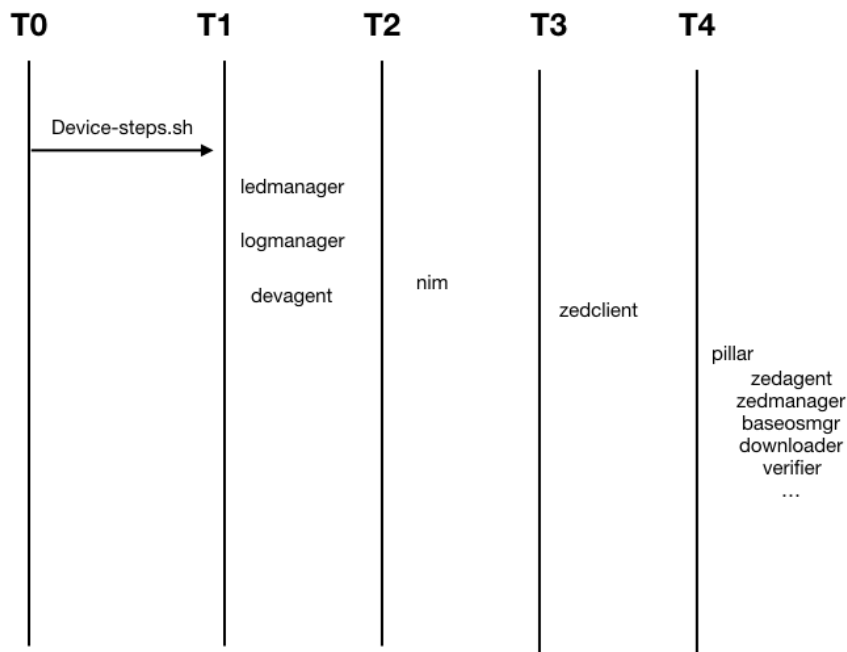
In the whole EVE node boot up process, ZedAgent and associated modules are spawned, only after network connectivity(through nim, waitfor address) and device registration (zedclient).

For baseos upgrade validation, this leaves a gap between node boot up and real baseos upgrade transition process invocation in zedagent. Any failure inbetween, the device boot up until zedagent starts, may lead to device being struck in some indefinite state and may turn the device to a non-functional unit.



Proposal

The zedagent module will be broken-up. The baseos upgrade validation and device health will be managed by NodeAgent. The NodeAgent will be one of the first modules to be spawned along with ledmanager, and will be persistent for the whole lifetime of the EVE node. The ZedAgent will be only responsible for cloud connectivity and configuration parsing and status/metrics publication. Baseosmgr will interact with NodeAgent for the baseos upgrade installation and validation.



EVE Node Health Monitor Function

EVE Node health check functionality, currently consists of the following,

Watchdog Time : For Pillar Agent(s) Health and responsiveness.

Each agent's health is monitored through software watchdog. If an agent does not retouch the pid file for watchdog time interval, the device is rebooted.

Reset Time: For controller connectivity health in normal operation mode

On controller connectivity loss, the EVE node is rebooted after the reset time interval.

Fallback Time: For controller connectivity health during baseos upgrade validation

For controller connectivity loss, EVE Node reboots and falls back to fallback image, after the fallback time interval.

Current Implementation

The watchdog time handler functionality is based on `wdctl` utility, and it is part of `device-steps.sh`.

The reset and fallback time functionalities are currently part of ZedAgent Module.

Refactoring Details

The watchdog time functionality will remain as such. The reset and fallback time functionality will be moved into a new agent called, NodeAgent. The whole baseos upgrade validation orchestration functionality will be moved into NodeAgent module. NodeAgent will be spawned along with ledmanager. NodeAgent will listen to ledmanager ledblinker config messages to determine controller connectivity status along with successful configuration pull message time stamps from zedagent, to orchestrate the baseos upgrade validation functionality. NodeAgent will be owner for Zboot config and will publish them for usage by baseosmanager. Also on successful baseos installation and reset/fallback timer expiry, the device reboot operations will be triggered through "NodeAgent status" pubsub topic.

Zedagent module will only be responsible for controller connectivity related functionalities, like pulling latest configuration blob from controller, and publishing status/info/metrics messages to controller. And will update this information through "zedagent status" pubsub topic. Zedagent will subscribe to "NodeAgent status" pubsub topic to execute device reboot commands.

Baseosmanger will listen to NodeAgent module, zboot config messages to handle, and update zboot status, for baseos installation and upgrade validation orchestration.

In a nutshell, the following are going to be changes in event handling per module.

Baseosmgr Module

Baseosmgr will subscribe to the following topic,

- "zboot config" from NodeAgent
For baseos installation and upgrade validation

ZedAgent Module

Zedagent will subscribe to the following topic,

- "NodeAgent status" , generated by NodeAgent
For executing device reboot command
To publish the remaining test time to controller, for baseos upgrade validation

Zedagent will publish the following topic,

- "zedagent status"
Time stamp for last successful configuration pull from controller

NodeAgent Module

NodeAgent module will subscribe to the following topics,

- "ledBlinker config", generated by zedclient/zedagent, etc
For EVE node registration, controller connectivity change events
- "zboot status", generated by baseosmgr
For baseos installation and upgrade validation orchestration
- "zedagent status", generated by zedagent
For the last successful config fetch time stamp, from controller

NodeAgent will publish the following topics,

- "zboot config"
Zboot partition information
- "NodeAgent status"
For device reboot event, in baseos installation and reset/fallback timer expiry
Remaining test time, for publication to controller (consumed by zedagent)

P.S.

For completeness and future workscope, the following items are noted, for EVE node health. This list is not exhaustive, and the necessary actions for them needs be defined.

- cpu usage health

- disk space usage health
- network usage health
- each agent's basic functionality check, (on upgrade)
- controller driven testing and marking the baseos as active