Ref. Ares(2016)5382708 - 15/09/2016

MONROE

Measuring Mobile Broadband Networks in Europe

H2020-ICT-11-2014
Project number: 644399

Deliverable D4.1

Maintenance Routines

Editor(s): Özgü Alay
Contributor(s): Thomas Hirsch, Audun Fosselie Hansen, Andra Lutu, Miguel Peon Quiros, Jonas Karlsson

Work Package: 4 / Maintenance
Revision: 0.1
Date: September 15, 2016
Deliverable type: DEM (Report)
Dissemination level: Confidential, only for members of the consortium (including the Commission Services)

European Commission | Horizon 2020 European Union funding for Research & Innovation

Abstract

This document presents a description of maintenance goals and the corresponding maintenance procedures both for software and hardware components in the MONROE system. The procedures we propose in this document should be considered as a starting point within the Maintenance Work Package (WP4). Therefore, we will maintain this document as a living document. We will upgrade the maintenance procedures, hence update this document, throughout the project's lifetime.

Participant organisation name	**Short name**
SIMULA RESEARCH LABORATORY AS *(Coordinator)*	SRL
CELERWAY COMMUNICATION AS	CWY
TELENOR ASA	Telenor
NEXTWORKS	NXW
FUNDACION IMDEA NETWORKS	IMDEA
KARLSTADS UNIVERSITET	KaU
POLITECNICO DI TORINO	POLITO

Contents

1 Introduction

The main goal of the MONROE project is to build and operate a unique platform for measurements and experiments in operational mobile networks. Maintenance of such a large-scale platform is a crucial task. Clearly, experimenters making use of the testbed depend on its availability for their experiments. MONROE will provide maintenance for the nodes and the backend system, as well as the measurement data. The maintenance activities involve both hardware (HW) and software (SW) maintenance, including the replacement of the nodes when necessary. For all maintenance activities, we created a Maintenance Ticketing system where we follow maintenance issues with an issue tracker, log maintenance actions and distill repeating actions into maintenance routines.

In this document, we first describe the maintenance goals. We then focus on the maintenance procedures for node, backend and measurement data, and the procedures to be established. Finally, we discuss the procedures for SW bugs and upgrades.

2 Maintenance Goals

The consortium has defined the three main maintenance goals as follows:

- **Keep alive:** Keep the nodes in operating condition (i.e. accessible and being available to experiments) and keep the number of nodes that are operating large enough at all times.
- **React fast:** Automate the maintenance procedure as much as possible in order to react quickly to known anomalies. In order to take action within 5 days from when a problem occurs, we check the status of the nodes at least twice a week. If the problem cannot be solved remotely, we replace the nodes within 2 weeks for stationary nodes and 4 weeks for mobile nodes.
- **Data first:** No measurement data should be lost once the data is transferred from the nodes to the back-end.

All maintenance events will be tracked in the Maintenance Ticketing system[1] in github repository.

3 Node Maintenance

Node maintenance considers both HW and SW maintenance for the nodes. The consortium will assign a representative to each country and each representative will check the node status twice a week. The report on the status of the nodes during the last week will be automated by sending an email summarizing this report to the maintenance team. Actions that need to be taken after this status check are defined below, separately for HW and SW.

3.1 Node HW Maintenance

We consider two main cases for node HW maintenance: (i) the node is completely lost or (ii) part of the node is operational. Below, we define the procedure for node HW maintenance in these two cases.

[1] https://github.com/MONROE-PROJECT/Maintenance

3.1.1 Node is unreachable

We define the HW maintenance procedures when the node is not reachable differently for stationary and mobile nodes.

Stationary Nodes: When a stationary node is **NOT** online for 24 hours:

- Power cycle the node by sending an SMS through the GSM socket
- If the node is still not online after 36 hours, *manual intervention* is required.

Mobile Nodes: When a mobile node is NOT online for 3 days:

- Check if bus/train/truck is still operating
- If bus/train/truck is still operating, *manual intervention* is required.

Manual Intervention: If the node cannot be recovered for a certain period of time, manual intervention is required. Our approach to manual intervention is to replace the node (with the test nodes). To this end, the representative will contact the host to arrange the replacement of the node. Once the node is in the lab, the engineering team will perform the troubleshooting and identify the components that failed the most. We will then buy backup parts for these components. In the beginning of the maintenance activities, backup components has been ordered to replace 10 nodes.

For the mobile nodes, since the installation requires more effort, before the replacement of the node, the engineer on the site will access to the node over the management ethernet cable and will try to see whether there is a SW issue that can be solved on site.

3.1.2 Node is partially available

In some cases the node can be reachable, but some components may not be operational. For example, some or all of the Mifi's can be unreachable, or internal management interface can be offline while Mifi's are operational. We define below the procedures to handle different cases where the node is partially available:

Mifi Maintenance:

If a Mifi is gray in the inventory for more than 1 hour while the node is online and there exists other Mifis that are online, the first action is to restart the Mifi. Policies for automatically restarting the Mifi (e.g. when no service) can be summarized as:

- Restart the USB port and/or Mifi
- If the Mifi is still not visible in lsusb, after 2 days, drain the battery of the Mifi to trigger hard restart of the Mifi.
- If the Mifi is still not available after we drain the battery, the Mifi has to be replaced.
  - For stationary nodes, replace the non-working Mifi only.
  - For mobile nodes, replace the whole node with a test node.

USB Hub Maintenance:
If all the Mifis are not reachable, the root cause of the problem can be the USB hub. In that case, the maintenance procedure is as follows:

- Check error message in logs and reboot node if the error messages indicate that the hub has a problem.
- If the reboot has no effect, replace the node.

WiFi and Sierra Wireless Maintenance:
Replace the node if WiFi and/or Sierra is not working for more than 14 days.

3.2 SW Maintenance

Software on the node consists of several layers, all of which have to be operational to make the node available for experimenters, and to ascertain correct measurement results.

Below we list all different software components that are running on the node:

- maintenance access
- software watchdog
- kernel patches
- core software (routing, load balancer, network-listener)
- sensors
- metadata-exporter
- monroe base experiments
- usage monitoring
- scheduling client
- container and virtualization system

Maintenance activity for software consists of reacting to software updates and bugs as well as resource exhaustion and availability issues in a timely manner. Policies regarding SW bugs and updates will be discussed in Section 5.1. Below, in indor to react on resource exhaustion and availability issues, we describe the policies for network assessment and automation via watchdogs. We have further implemented a maintenance mode and maintenance window as detailed below.

Network assessment: Very low data throughput checks will be run every 5 minutes in order to assess all different interfaces. The management interface is default for all maintenance activities, however, if the management interface is not available due to a problem or coverage, the node will switch to the mifi interface with the best quality. This guarantees that the maintenance activities can be carried out at all times.

Hardware Watchdogs for rebooting the node: The hardware watchdog is a hardware component that triggers a reboot if the system becomes unresponsive, e.g. in the case of a kernel failure. The watchdog monitors whether a device file has been written to at least once. If it has been written to, but no subsequent writes occur within a defined period, a reboot is triggered. A software daemon is running on the node which writes to the device in the given interval, preventing the node from rebooting. This should always be the case, except in the case of severe system corruption.

Software Watchdogs for rebooting the node: The daemon[2] which writes to the watchdog device can run scripts or other binaries defining a test procedure, and a repair procedure to run if the test fails. Should a repair action fail, the watchdog will trigger a reboot. Thus, we can trigger a repair action and/or a reboot through arbitrary conditions. This watchdog is thus designed for monitoring system level operations, more complex tests are handled by the monroe watchdog below.

MONROE Watchdog: Another, higher-level watchdog is run in regular intervals to monitor all necessary services. This watchdog can implement more complex tests, repair actions, and resolutions. In particular, it can decide to set the node in Maintenance mode (below), trigger a reboot, or initiate a complete reinstallation of the base system through the boot OS. Only if the tests defined in this watchdog succeed, a successful boot flag is written. Repeated failures to boot the node into a working state will eventually trigger a reinstallation of the system by the boot OS.

Maintenance mode: When the node is up but configuration errors are observed or certain tests fail, the node will go into maintenance mode. In this mode, all docker containers will be stopped immediately and a message will be written to syslog (every minute). Also in the scheduler, the node status will be set to maintenance and no new experiment will be allowed. An email will be sent to the representative of the node and troubleshooting will be initiated. Maintenance mode can only be reversed by the engineering team once the issue is resolved.

Node Maintenance Hour: Every day, all stationary nodes and their Mifis are restarted. This maintenance window can also be used for SW updates and status checks.

3.3 What can be automated?

In order to run the maintenance activities efficiently and to react timely, we aim to automate as many tasks as possible. Below we provide list of items that will be automated and this list will be updated throughout the course of the project:

- Soft reboot the node when
  - SW modules not working properly
  - the node is without network (i.e. all the interfaces are down)
  - certain error messages are seen in the log files
- Reboot a MiFi when
  - a MiFi is without network

[2] http://linux.die.net/man/8/watchdog

  - a Mifi has lost the communication with the node

- Drain the Mifi battery when we cannot communicate with the MiFi and if soft reboot did not work
- Restart of sierra wireless when it is not connected to the network
- Restart of WiFi when it is not connected to the network

4 Backend and DATA Maintenance

Backend and DATA maintenance considers both HW and SW maintenance for continuously running the different backend services. As with the software running on the node, these services are crucial to ensure the measurement results are generated, transferred and stored correctly.

The complete list of services is as follows:

- Cassandra Database
- Database Importer
- Scheduling server
- Result storage
- User Access
- Measurement responder
- Visualization
- NTP
- Certification Authority / User accounts

The consortium will assign a representative for each service (or a group of services) and the representatives will check the status of theses services twice a week. The report on the status of maintenance activities will be shared among the representatives.

Below we define the maintenance activities separately for HW, SW and data maintenance. Most of these services take the form of a web server running dedicated software, hence maintenance procedures are similar and aligned with commonplace server operations.

4.1 HW Maintenance

- Check and ensure that the servers are powered, running, and have resources to operate for the next month.
- Monitor disk space, CPU load, memory, etc... and upgrade when necessary.

4.2 SW Maintenance

- Check and ensure that the services are running, and monitored by a software watchdog to be restarted on failure.
- Check and ensure that software running on the server is up to date and receives necessary security updates.
- Check that the components that depend on this service are able to communicate and retrieve the correct results from the service.
- Backup of the SW: All software and configuration is stored in a public or a private git repository on the MONROE github and bitbucket accounts. Compiled software packages are backed up as well.

4.3 Data Maintenance

Data maintenance encompasses backups and validation of:

- Measurement results of MONROE base experiments.
  - Source data (node json files): Automatically compressed and backed-up every night to a remote location (IMDEA).
  - Online DB replicas: The MONROE database will be replicated to two or more servers located at partners within the project (depending on SW/HW availability).
  - Additional offline DB: Kept (e.g., at IMDEA) for direct access by the external experimenters; this DB instance will be updated daily from the log files and will remain disconnected from the main copies.
- System log files from the nodes and backend servers are sent to a central SYSLOG server.
- Result files and container logs from user experiments are transferred to a central server for later retrieval by the experimenters. These files will be kept for two weeks; then, they will be deleted.

Additional processes for data maintenance include:

- Ensuring that the data is monitored and backed up regularly, including the log files.
- A service will run in the database server to monitor DB activity and alert the corresponding representative if no data has been inserted for any specific node in the (online) DB during the last 30 minutes so that manual procedures can be started. Additional validation procedures for the last received data may be performed by this service.
- Restarting the insertion process of metadata files into the database on failure.
- Deciding when to delete the log files (if space issues arise).

5 Procedures for SW Bugs and Upgrades

5.1 SW bugs

Below, we define the procedure we follow when we discover a SW bug:

- Bugs are rated as
  - Critical: may affect node stability
  - Important: may affect experiment results
  - Other: minor SW bugs
- Critical and important bugs are published on the open call mailing list and MONROE developer mailing list.
- Fixes for critical bugs are pushed immediately, after following the testing procedure. If a downtime is necessary, this is published on the open call mailing list and monroe developer mailing list.
- Fixes for important bugs are published in the next update window (the frequency of the update window will be adjusted during the lifetime of the project).
- Other bugs are published in the update windows.

5.2 SW upgrades

Below, we define the procedure for SW upgrades, both on the nodes and in the backend system:

- Updates are pushed during the maintenance hour and should install within 60 minutes.
- Non-critical updates are pushed once per week, and an update report is written to the development mailing list.
- Critical updates are pushed during the next maintenance window, and an update report written to the development mailing list.
- Configuration and package management follows the separation development/staging/production. Dedicated staging nodes will run the updated configuration for at least 7 days before pushing a non-critical update.

Disclaimer

The views expressed in this document are solely those of the author(s). The European Commission is not responsible for any use that may be made of the information it contains.

All information in this document is provided "as is", and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.