
FIWARE Monitoring

Release

January 13, 2016

1	Introduction	1
1.1	FIWARE Monitoring	1
1.2	User and Programmers Guide	6
1.3	Installation and Administration Guide	9

Introduction

Monitoring GE - TID Implementation is the key component to allow incorporating monitoring and metering mechanisms in order to be able to constantly check the performance of the system, but the architecture should be easily extended to collect data for other required needs. Monitoring involves gathering operational data in a running system. Collected information can be used for several purposes:

- Cloud users to track the performance of their own instances.
- SLA management, in order to check adherence to agreement terms.
- Optimization of virtual machines.

The monitoring system is used by different Cloud GEs in order to track the status of the resources. They use gathered data to take decisions about elasticity or for SLA management. Whenever a new resource is deployed in the cloud, the proper monitoring probe is set up and configured to start providing monitoring data.

The FIWARE Monitoring source code can be found [here](#)

This documentation offers deeper information on FIWARE Monitoring.

Documentation

1.1 FIWARE Monitoring

This is the code repository for FIWARE Monitoring, the reference implementation of the Monitoring GE.

This project is part of [FIWARE](#). Check also the [FIWARE Catalogue entry for Monitoring](#).

Any feedback on this documentation is highly welcome, including bugs, typos or things you think should be included but aren't. You can use [github issues](#) to provide feedback.

For documentation previous to release 4.4.2 please check the manuals at FIWARE public wiki:

- [FIWARE Monitoring - Installation and Administration Guide](#)
- [FIWARE Monitoring - User and Programmers Guide](#)

1.1.1 GEi overall description

FIWARE Monitoring is the key component to allow incorporating monitoring and metering mechanisms in order to be able to constantly check the performance of the cloud infrastructure.

This involves gathering operational data in a running system, which usually requires collecting data from heterogeneous sources. Besides, the monitoring architecture should be easily extended to collect additional data for any other required needs.

FIWARE Monitoring is agnostic to the *framework* used to gather monitoring data. It just assumes there are several *monitoring probes* collecting information, which somehow must be forwarded to an *adaptation layer*, responsible for transforming data into a common representation (*NGSI*) and publishing through a *Context Broker* (see [Orion](#)).

Collected information can be used for several purposes:

- Cloud users to track the performance of their own instances.
- SLA management, in order to check adherence to agreement terms.
- Optimization of virtual machines.

Components

Monitoring framework It is up to the infrastructure owner which tool (like [Nagios](#), [Zabbix](#), [openNMS](#), [perfSONAR](#), etc.) is installed for this purpose.

Collector Framework-specific component to forward monitoring data being gathered to the adaptation layer (i.e. *NGSI Adapter*). Monitoring GE provides a Nagios loadable module NGSI Event Broker as collector for such monitoring framework.

Adaptation layer NGSI Adapter serves as generic adapter to transform monitoring data from probes to NGSI context attributes.

1.1.2 Build and Install

The recommended procedure is to install using RPM packages in CentOS 6.x, or DEB packages in Ubuntu 12.04/14.04 LTS. If you are interested in building from sources, check this document.

Requirements

- System resources: see these recommendations.
- Operating systems: CentOS (RedHat) and Ubuntu (Debian), being CentOS 6.3 the reference operating system.
- RPM/DEB dependencies: some required packages may not be present in official repositories, or their versions are too old (for example, `nodejs`). In any case, checking for such dependencies and configuration of alternative sources is automatically managed by the package installation scripts when using the proper tool (`yum` in CentOS or `apt-get/gdebi` in Ubuntu).

Installation

Using FIWARE package repository (recommended)

Refer to the documentation of your Linux distribution to set up the URL of the repository where FIWARE packages are available (and update cache, if needed):

CentOS

`http://repositories.testbed.fiware.org/repo/rpm/x86_64`

Ubuntu

```
http://repositories.testbed.fiware.org/repo/deb
```

Then, use the proper tool to install the packages (this depends on monitoring framework used in the cloud infrastructure, but at least NGSI Adapter will be installed in any case):

CentOS

```
$ sudo yum install fiware-monitoring-ngsi-adapter
```

Ubuntu

```
$ sudo apt-get install fiware-monitoring-ngsi-adapter
```

Additionally, in case Nagios 3.4/3.5 and its probes (*Nagios Plugins*) are being used as the framework to gather monitoring data, then we may install the package `fiware-monitoring-ngsi-event-broker` (see [Components](#) above).

Using the RPM/DEB files

Download the package(s) from the [FIWARE Files area](#) and use the proper tool to install it. Take into account that you may need to manually install dependencies, as some tools aren't able to manage them when installing from file:

CentOS

```
$ sudo rpm -i fiware-monitoring-ngsi-adapter-X.Y.Z-1.noarch.rpm
$ sudo rpm -i fiware-monitoring-ngsi-event-broker-X.Y.Z-1.x86_64.rpm
```

Ubuntu

```
$ sudo dpkg -i fiware-monitoring-ngsi-adapter_X.Y.Z_all.deb
$ sudo dpkg -i fiware-monitoring-ngsi-event-broker_X.Y.Z_amd64.deb
```

Upgrading from a previous version

Unless explicitly stated, no migration steps are required to upgrade to a newer version of the Monitoring components:

- When using the package repositories, just follow the same directions described in the [Installation](#) section (the `install` subcommand also performs upgrades).
- When upgrading from downloaded package files, use `rpm -U` in CentOS, or use same `dpkg -i` command in Ubuntu.

1.1.3 Running

As explained in the [overall description](#) section, there are a variety of elements involved in the monitoring architecture, apart from those components provided by this Monitoring GE (at least, an instance of *Context Broker* is required and some underlying monitoring framework, such as *Nagios*). Please refer to their respective documentation for instructions to run them. From the Monitoring GE components, only NGSI Adapter runs as standalone server. Once installed, there are two ways of running NGSI Adapter: manually from the command line or as a system service (the latter only available if installed as a package). It is not recommended to mix both ways (e.g. start it manually but use the service scripts to stop it). This section assumes you are using the system service (recommended): for the command line alternative, please refer to this document.

In order to start the adapter service, run:

```
$ sudo service ngsi_adapter start
```

Then, to stop the service, run:

```
$ sudo service ngsi_adapter stop
```

We can also force a service restart:

```
$ sudo service ngsi_adapter restart
```

Configuration file

The configuration used by the adapter service is optionally read from the file `/etc/sysconfig/ngsi_adapter` (in CentOS) or `/etc/default/ngsi_adapter` (in Ubuntu):

```
# ADAPTER_LOGFILE - Logging file
ADAPTER_LOGFILE=/var/log/ngsi_adapter/ngsi_adapter.log

# ADAPTER_LOGLEVEL - Logging level
ADAPTER_LOGLEVEL=INFO

# ADAPTER_LISTEN_HOST - The host where NGSI Adapter listens to requests
ADAPTER_LISTEN_HOST=0.0.0.0

# ADAPTER_LISTEN_PORT - The port where NGSI Adapter listens to requests
ADAPTER_LISTEN_PORT=1337

# ADAPTER_BROKER_URL - The endpoint where Context Broker is listening
ADAPTER_BROKER_URL=http://127.0.0.1:1026/

# ADAPTER_RETRIES - Maximum number of retries invoking Context Broker
ADAPTER_RETRIES=2
```

Most of these attributes map to options of the command line interface as follows:

- `ADAPTER_LOGLEVEL` maps to `-l` or `--logLevel` option
- `ADAPTER_LISTEN_HOST` maps to `-H` or `--listenHost` option
- `ADAPTER_LISTEN_PORT` maps to `-p` or `--listenPort` option
- `ADAPTER_BROKER_URL` maps to `-b` or `--brokerUrl` option
- `ADAPTER_RETRIES` maps to `-r` or `--retries` option

Default values are found in `/opt/fiware/ngsi_adapter/config/options.js`.

Checking status

In order to check the status of the adapter service, use the following command (no special privileges required):

```
$ service ngsi_adapter status
```

1.1.4 API Overview

To transform monitoring data into NGSI attributes, probe raw data should be sent as body of a POST request to the adapter, identifying the source entity being monitored in the query fields.

For example, if using the `check_load` Nagios probe to measure CPU load, then the request would look like:

```
curl "{adapter_endpoint}/check_load?id={myhostname}&type=host" -s -S --header 'Content-Type: text/plain'
OK - load average: 5.00, 7.01, 7.05|load1=5.000;10.000;10.000;0; load5=7.010;15.000;15.000;0; load15=7.050;30.000;30.000;0;
EOF
```

This would result in an invocation to Context Broker updating the context of an entity of type `host` identified by `myhostname` with a new attribute `cpuLoadPct` with value `5.00`.

Please have a look at the [API Reference Documentation](#) section below and at the programmer guide.

API Reference Documentation

- [FIWARE Monitoring v1 \(Apiary\)](#)

1.1.5 Testing

End-to-end tests

Please refer to the Installation and administration guide for details.

Unit tests

The `test` target is used for running the unit tests in both components of Monitoring GE:

```
$ cd ngsi_adapter
$ grunt test

$ cd ngsi_event_broker
$ make test # synonym of standard 'check' target
```

Please have a look at the section building from source code in order to get more information about how to prepare the environment to run the unit tests.

Acceptance tests

In the following documents you will find a business readable description of the features provided by the components of the Monitoring GE, as well as automated tests for them:

- [NGSI Adapter acceptance tests](#)

1.1.6 Advanced topics

- Installation and administration
 - Building from sources
 - Running Adapter from command line
 - Logs
 - Resources & I/O Flows
- User and programmers guide
 - NGSI Adapter custom probe parsers

- Retrieval of historical data

1.1.7 License

(c) 2013-2015 Telefónica I+D, Apache License 2.0

1.2 User and Programmers Guide

1.2.1 Introduction

Welcome the User and Programmers Guide for the Monitoring Generic Enabler. This GE is built up from different distributed components, as depicted in the following figure:

Background and Detail

This User and Programmers Guide relates to the Scalability Manager GE which is part of the [Cloud Hosting Chapter](#). Please find more information about this Generic Enabler in the following [Open Specification](#).

1.2.2 User Guide

This GE does not provide an interactive user interface, hence there is no User Guide. The following section elaborates on programmatic usage.

1.2.3 Programmer Guide

According to the architecture aforementioned, there are several APIs involved in the monitoring process:

- NGSI Adapter API
- Context Broker API
- Monitoring (*Query Manager*) API

NGSI Adapter API

Probe raw data should be sent as body of a POST request to the adapter, identifying the source entity being monitored in the query parameters. For example, given the following scenario:

Monitored host 178.23.5.23

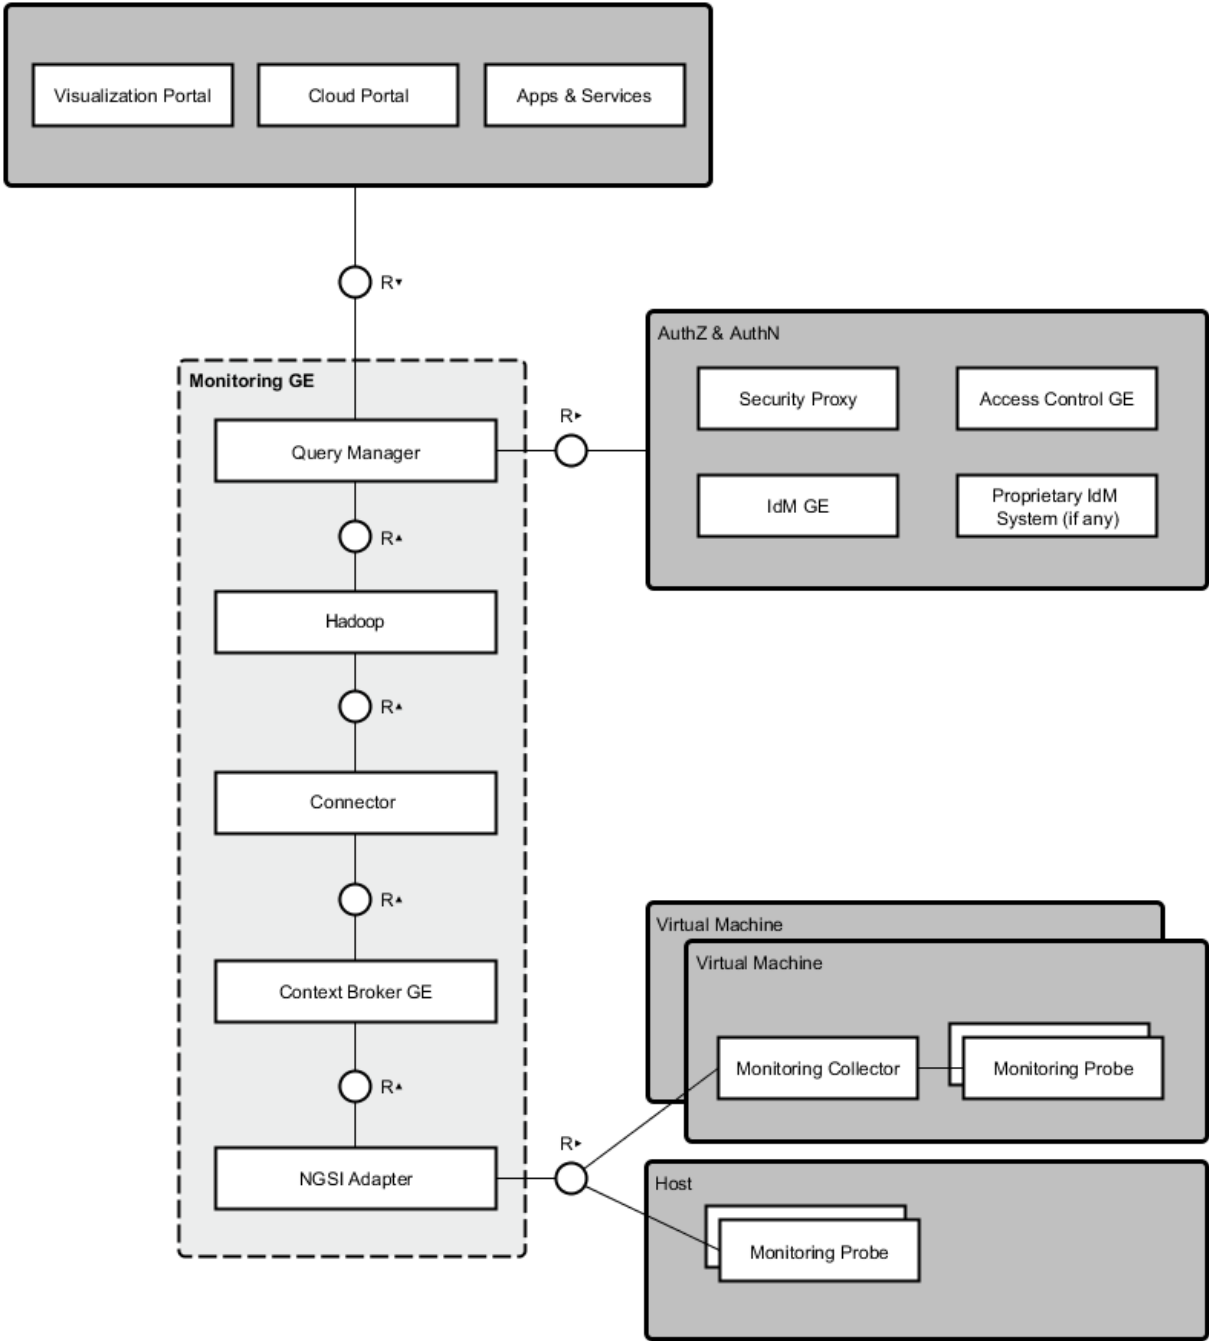
Monitoring tool Nagios

Monitoring probe name check_load

NGSI Adapter endpoint http://adapterhost:1337

then requests would look like:

```
HTTP POST http://adapterhost:1337/check_load?id=178.23.5.23&type=host
Content-Type: text/plain
OK - load average: 0.36, 0.25, 0.24|load1=0.360;1.000;1.000;0; load5=0.250;5.000;5.000;0; load15=0.240;1.000;1.000;0;
```



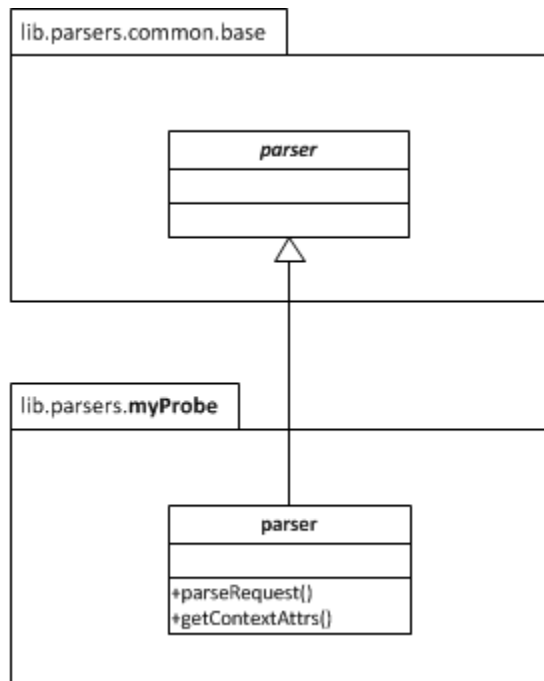
Please take into account that NGSI standard identify entities (in this case, the resources being monitored) using a pair `<entityId,entityType>`. This identification of the monitored resource has to be provided as the query parameters `id` and `type`, respectively. The probe name included in the URL lets NGSI Adapter know the originating monitoring probe, therefore selecting the proper parser for it. This API is fully described in [Apiary](#).

Monitoring framework is expected to schedule the execution of probes and send the raw data been gathered to the NGSI Adapter. Depending on the tool that has been chosen, this would require the development of a custom component (a kind of **monitoring collector**) to automatically forward such data to the adaptation layer.

NGSI Adapter parsers

NGSI Adapter processes requests asynchronously, trying to locate a valid parser named after the originating probe, located at `lib/parsers/`. If probe is unknown, response status will be 404; otherwise, response status will be 200, parser will be dynamically loaded, and then its `parseRequest()` and `getContextAttrs()` methods will be invoked. With the attribute list returned by the latter, Context Broker will be invoked.

Custom parsers for new probes may be easily added to NGSI Adapter, just extending a base abstract object and implementing the aforementioned methods. For example, suppose we want to support a new “*myProbe*” whose data is a comma-separated list of values of two attributes *myAttr0* and *myAttr1*:



```
// module "lib.parsers.myProbe"

var baseParser = require('./common/base').parser,
    myParser    = Object.create(baseParser);

myParser.parseRequest = function(request) {
  return { data: request.body };
};

myParser.getContextAttrs = function(probeData) {
  var items = probeData.data.split(',');
  return { myAttr0: items[0], myAttr1: items[1] };
};
```

```
};  
  
exports.parser = myParser;
```

Context Broker API

Please refer to [Context Broker documentation](#). This will give us access to the last updates of monitoring data available, but not to historical data.

Monitoring API

Retrieval of historical data stored at a distributed filesystem (e.g. Hadoop) is handled by the *Query Manager* component, whose API is described in this [preliminary specification](#).

1.3 Installation and Administration Guide

1.3.1 Introduction

This guide defines the procedure to install the different components that build up the Monitoring GE, including its requirements and possible troubleshooting.

For general information, please refer to this document.

1.3.2 Installation

Monitoring infrastructure comprises several elements distributed across different hosts, as depicted in the following figure:

Installation of probes


Monitoring GE is agnostic to the framework used to gather monitoring data. It just assumes there are several probes collecting such data, which somehow will be forwarded to the adaptation layer (NGSI Adapter).

It is up to the infrastructure owner which tool (like [Nagios](#), [Zabbix](#), [openNMS](#), [perfSONAR](#), etc.) is installed for this purpose.

Installation of collector

Probes must “publish” their data to NGSI Adapter. Depending on the exact monitoring tool installed, a kind of *collector* has to be deployed in order to send data to the adapter:

- **NGSI Event Broker** is an example specific for Nagios, implemented as a loadable module. Description and installation details can be found [here](#).



1. **Probes** gather raw monitoring data, which a **Collector** (for Nagios, this is *NGSI Event Broker*) forwards to *NGSI Adapter*.
2. **NGSI Adapter**, responsible for translating probe raw data into a common format (NGSI).
3. **Parsers** at NGSI Adapter, specific for the different probes that generate monitoring data.
4. **Context Broker**, where monitoring data (transformed into NGSI context updates) will be published.
5. **Hadoop**, for storing historical context data.
6. **Connector** between Context Broker and data storage (for example, this could be *Cygnus*).

Installation of NGSI Adapter

Requirements

NGSI Adapter should work on a variety of operating systems, particularly on the majority of GNU/Linux distributions (e.g. Debian, Ubuntu, CentOS), as it only requires a V8 JavaScript Engine to run a Node.js server.

Hardware Requirements The minimal requirements are:

- RAM: 2 GB

Software Requirements NGSI Adapter is a standalone Node.js process, so `node` and its package manager `npm` should be installed previously. These requirements are automatically checked when installing the `fiware-monitoring-ngsi-adapter` package. However, for manual installation please visit [NodeSource](https://nodejs.org/en/).

Downloads

Please refer to this document for details.

Additional parsers

NGSI Adapter currently includes a predefined set of parsers for Nagios probes at `lib/parsers/` directory, each parser named after its corresponding probe.

Whenever new Nagios probes are used, or even if we decide to use a different monitoring framework with its own set of probes, the corresponding parsers should be dropped in such directory.

Installation of Context Broker

Please refer to [Orion](#) documentation.

Installation of the connector

This component subscribes to changes at Context Broker and writes data into a distributed filesystem storage (usually HDFS from [Hadoop](#)). Historically the **ngsi2cosmos** connector implementation has been used (installation details [here](#)), although from March 2014 this component is deprecated and a brand new **Cygnus** implementation (installation details [here](#)) is available.

1.3.3 Running the monitoring components

As stated before, there are a number of distributed components involved in the monitoring. Please refer to their respective installation manuals for execution details (this applies to probes & monitoring software, Context Broker, Hadoop, etc.). This section focuses on NGSI Adapter specific instructions.

Running NGSI Adapter

Once installed, there are two ways of running NGSI Adapter: manually from the command line or as a system service. It is not recommended to mix both ways (e.g. start it manually but using the service scripts to stop it).

As system service

When installed from its package distribution, a Linux service `ngsi_adapter` is configured (but not started). Please refer to this document for details.

From the command line

You can run the adapter just typing the following command at the installation directory (usually `/opt/fiware/ngsi_adapter/`):

```
$ adapter
```

You can use these command line options (available typing `adapter --help`):

-l, --logLevel	Verbosity of log messages
-H, --listenHost	The hostname or address at which NGSI Adapter listens
-p, --listenPort	The port number at which NGSI Adapter listens
-b, --brokerUrl	The URL of the Context Broker instance to publish data to
-r, --retries	Number of times a request to Context Broker is retried, in case of error

1.3.4 Sanity check procedures

These are the steps that a System Administrator will take to verify that an installation is ready to be tested. This is therefore a preliminary set of tests to ensure that obvious or basic malfunctioning is fixed before proceeding to unit tests, integration tests and user validation.

End to end testing

Use the commands of the monitoring framework being used (for example, Nagios) to reschedule some probe execution and force the generation of new monitoring data:

- Check the logs of the framework (i.e. `/var/log/nagios/nagios.log`) for a new probe execution detected by the *collector*:

```
$ cat /var/log/nagios/nagios.log
[1439283831] lvl=INFO | trans=rdPmJ/uHE62a | comp=fiware-monitoring-ngsi-event-broker | op=NGSI
```

- Check NGSI Adapter logs for incoming requests with raw data, and for the corresponding `updateContext()` request to Context Broker:

```
$ cat /var/log/ngsi_adapter/ngsi_adapter.log
time=... | lvl=INFO | trans=rdPmJ/uHE62a | op=POST | msg=Request on resource /check_xxx with par
time=... | lvl=INFO | trans=rdPmJ/uHE62a | op=POST | msg=Response status 200 OK
time=... | lvl=INFO | trans=rdPmJ/uHE62a | op=UpdateContext | msg=Request to ContextBroker at ht
```

- Finally, query Context Broker API to check whether entity attributes have been updated according to the new monitoring data (see details [here](#))

List of Running Processes

A node process running the “adapter” server should be up and running, e.g.:

```
$ ps -C node -f | grep adapter
fiware    21930      1  0 Mar28 ?           00:06:06 node /opt/fiware/ngsi_adapter/src/adapter
```

Alternatively, we can check if service is running, e.g.:

```
$ service ngsi_adapter status
* ngsi_adapter is running
```

Network interfaces Up & Open

NGSI Adapter uses TCP 1337 as default port, although it can be changed using the `--listenPort` command line option.

Databases

This component does not persist any data, and no database engine is needed.

1.3.5 Diagnosis Procedures

The Diagnosis Procedures are the first steps that a System Administrator will take to locate the source of an error in a GE. Once the nature of the error is identified with these tests, the system admin will very often have to resort to more concrete and specific testing to pinpoint the exact point of error and a possible solution. Such specific testing is out of the scope of this section.

Resource availability

Although we haven't done yet a precise profiling on NGSI Adapter, tests done in our development and testing environment show that a host with 2 CPU cores and 4 GB RAM is fine to run server.

Remote service access

- Probes at monitored hosts should have access to NGSI Adapter listen port (TCP 1337, by default)
- NGSI Adapter should have access to Context Broker listen port (TCP 1026, by default)
- Connector should have access to Context Broker listen port in order to subscribe to context changes
- Context Broker should have access to Connector callback port to notify changes

Resource consumption

No issues related to resources consumption have been detected neither with the NGSI Adapter server nor with the NGSI Event Broker loaded as a “pluggable” module on Nagios startup.

I/O flows

Figure at [installation section](#) shows the I/O flows among the different monitoring components:

- Probes send requests to NGSI Adapter with raw monitoring data, by means of a custom *collector* component (for example, NGSI Event Broker)
- NGSI Adapter sends request to Context Broker in terms of context updates of the monitored resources
- Context Broker notifies Connector with every context change
- Connector writes changes to storage