
CONFIGURING AN RVI NODE

This document describes the process of configuring an RVI node so that it can serve locally connected services, and also find other RVI nodes in a network.

READER ASSUMPTIONS

The reader is assumed to be able to:

1. Have a basic understanding of Linux directory structures.
2. Start and stop programs on the RVI-hosting system
3. Edit configuration files.
4. Understand the basic concepts of IP addresses, ports and URLs.

PREREQUISITES

1. Erlang runtime R16B01 or later has to be installed on the hosting system.
2. The `setup_rvi_node.sh` tool is available to build a release.
3. (Recommended) `priv/vehicle.config` is used as a starting point for a customized setup. Root access is not needed.

CONFIGURATION PROCESS OVERVIEW

To bring up an RVI node so that it can be used by locally connected services and communicate with other RVI nodes, the following steps must be taken.

1. Specify the node service prefix

This node will handle traffic to all services that start with the given prefix.

2. Specify RVI node external address

The external address is announced by the Data Link component to other RVI nodes, allowing them to connect to this node and exchange services.

2. Configure static nodes

Backend / Cloud-based RVI nodes have non-changing network addresses that should be known by other nodes in a network. This is achieved by setting up service prefixes and addresses of the static nodes in all other nodes

deployed in a network.

3. Specify Service Edge URL that local services connect to

The Service Edge URL is used by local services to send traffic that is to be forwarded to services on the local and remote nodes.

4. Specify URLs for RVI components

In addition to the Service Edge URL, the remaining components must have their URLs configured so that the components can locate each other and exchange commands.

5. Build the development release

The `setup_rvi_node.sh` is executed to read the configuration file and generate a development or production release.

6. Start the release

The `rvi_node.sh` is executed to launch the built development release.

CONFIGURATION FILE LOCATION AND FORMATS

There is a single configuration file, with the setup for all components and modules in the node, used for each release. All files are stored in the `priv` directory. A documented example file is provided as `rvi_sample.config`

The configuration file consists of an array of erlang tuples (records / structs / entries), where the `env` tuple contains configuration data for all components. The `rvi` tuple under `env` has all the configuration data for the RVI system. With the possible exception for the lager logging system, only the `rvi` tuple needs to be edited.

The term tuple and entry will be intermixed throughout this document.

SPECIFY NODE SERVICE PREFIX

All RVI nodes with locally connected services will announce these toward other, external RCI as a part of the service discovery mechanism. When announcing its local services to external RVI nodes, a node will prefix each service with a static string that is system-wide unique.

When a service sends traffic to another service, the local RVI node will prefix match the name of the destination service against the service prefix of all known nodes in the system. The node with the longest matching prefix will receive the traffic in order to have it forwarded to the targeted service that is connected to it.

The prefix always starts with an organisational ID that identifies the entity that manages the service. Best practises is to use the domain name of the hosting organisation.

Since every node's service prefix must be unique, they often contain a network address, a device id, a phone number, or similar device-unique information. Backend / Cloud nodes often have a symbolic, and unique prefix identifying what their role is.

Below are a few examples of prefixes:

`jaguarlandrover.com/vin/sajwa71b37sh1839/` - A JLR vehicle with the given vin.

`jaguarlandrover.com/mobile/+19492947872/` - A mobile device with a given number, managed by JLR, hosting an RVI node.

`jaguarlandrover.com/sota/` - JLR's global software over the air server.

`jaguarlandrover.com/3rd_party/` - JLR's 3rd party application portal.

`jaguarlandrover.com/diagnostic/` - JLR's diagnostic server.

The prefix for an RVI node is set in the `node_service_prefix` tuple.

An example entry is given below:

```
[
  ...
  { env, [
    ...
    { rvi, [
      ...
      { node_service_prefix, "jaguarlandrover.com/backend/" }
    ]}
  ]}
]
```

SPECIFY RVI NODE EXTERNAL ADDRESS

The external rvi node address is the address, as seen from the outside world, where this node's data link can be contacted. In IP based networks, this is usually a `hostname:port` value. In SMS-only networks, this will be the MSISDN of the node's mobile subscription. Any traffic directed to the given address should be forwarded to the Data Link component.

The configuration element to set under the `rvi` tuple is `node_address`.

An example tuple is given below:

```
[
  ...
  { env, [
    ...
    { rvi, [
      ...
      { node_address, "92.52.72.132:9850" }
    ]}
  ]}
]
```

Please note that IP addresses, not DNS names, should be used in all network addresses.

In the default Data Link component, `data_link_bert_rpc`, you also need to specify the port it should listen to, and possibly also the interface to use.

This is done by editing the tuple `rvi -> data_link -> bert_rpc_server`, and set `port` to the port that traffic is received on. If `data_link_bert_rpc` is to listen for traffic on only one interface, the IP address can be specified as `ip`.

An example tuple is given below:

```
[
  ...
  { env, [
    ...
    { rvi, [
      ...
      { components, [
        ...
        { data_link, [
          ...
          { bert_rpc_server, [
            { port, 9850 },
            { ip, "192.168.11.234" }
          ]}
        ]}
      ]}
    ]}
  ]}
]
```

If `data_link_bert_rpc` is to listen to the port on all network interfaces, the `ip` tuple can be omitted.

CONFIGURE STATIC NODES

Some RVI nodes in a network, such as central backend servers, will always be available at static network addresses. Services on these static nodes should be made available to all other nodes in a network (given that network connectivity is available).

The service prefixes and network addresses of static nodes can be configured in all other nodes, making the static nodes globally available outside the regular, peer-to-peer service discovery mechanism.

When traffic targeting a remote service is received by the RVI node from a locally connected service, it will first try to locate the remote node hosting the destination service through the service discovery database. If this fails the statically configured nodes are searched, prefix matching the name of the destination service against the specified static nodes' service prefixes.

If there is a match, the request will be sent to the network address of the matching node. If there are multiple

matches the static node with the longest matching prefix will receive the traffic.

Static nodes are configured as a list of tuples under the `static_nodes` tuple.

An example entry is given below:

```
[
  ...
  { env, [
    ...
    { rvi, [
      ...
      { static_nodes, [
        { "jaguarlandrover.com/sota/", "92.52.72.132:9850" },
        { "jaguarlandrover.com/remote_diagnostic/", "92.52.72.132:9851" }
      ]}
    ]}
  ]}
]
```

Please note that IP addresses, not DNS names, should be used in all network addresses.

Please note that all nodes configured in `static_nodes` must be up and running, listening on their specified addresses, before the node with the `static_nodes` configuration entry comes up. All static nodes are connected to during the startup sequence, and if one or more are not available, the startup sequence will fail. This will be fixed in future versions.

SPECIFY SERVICE EDGE URL

The Service Edge URL is that which will be used by locally connected services to interact, through JSON-RPC, with the RVI node.

Other components in the RVI node use the same URL to send internal traffic to Service Edge.

The URL of Service Edge is specified through the `service_edge` tuple's `url` entry, read by the other components in the node to locate it. When a URL is specified for Service Edge, the port that it is to listen to must be synchronized as well, using the `exo_http_opts` tuple.

An example entry is given below:

```
[
  ...
  { env, [
    ...
    { rvi, [
      ...
      { components, [
        ...
      ]}
    ]}
  ]}
]
```

```

    { service_edge, [
      { url, "http://127.0.0.1:8811" },
      { exo_http_opts, [ { port, 8811 } ] }
    ]}
  ]}
]}
]]
]

```

Please note that IP addresses, not DNS names, should be used in all network addresses.

SPECIFY URLS OF REMAINING RVI COMPONENTS

The remaining nodes in an RVI system needs to have their URLs and listening ports setup as well. It is recommended that consecutive ports after that used for `service_edge` are used.

Below is an example of a complete port/url configuration for all components, including the `bert_rpc_server` entry described in the external node address chapter:

```

[
  ...
  { env, [
    ...
    { rvi, [
      ...
      { components, [
        ...
        { service_edge, [
          { url, "http://127.0.0.1:8811" },
          { exo_http_opts, [ { port, 8811 } ] }
        ]},
        { service_discovery, [
          { url, "http://127.0.0.1:8812" },
          { exo_http_opts, [ { port, 8812 } ] }
        ]},
        { schedule, [
          { url, "http://127.0.0.1:8813" },
          { exo_http_opts, [ { port, 8813 } ] }
        ]},
        { authorize, [
          { url, "http://127.0.0.1:8814" },
          { exo_http_opts, [ { port, 8814 } ] }
        ]},
        { protocol, [
          { url, "http://127.0.0.1:8815" },
          { exo_http_opts, [ { port, 8815 } ] }
        ]},
        { data_link, [

```

```

    { url, "http://127.0.0.1:8816" },
    { exo_http_opts, [ { port, 8816 } ] },
    { bert_rpc_server, [ {port, 8817 } ] }
  ]}
]}
]}
]}
]

```

Please note that IP addresses, not DNS names, should be used in all network addresses.

RUNNING MULTIPLE NODES ON A HOST

Multiple RVI nodes can be run simultaneously on a single host as long as their configured URLs and ports do not interfere with each other. The data link external

In the example below a second URL/port setup is shown, using port range 9011-9017, that can co-exist with the setup listed in the examples in the previous chapters.

```

[
  ...
  { env, [
    ...
    { rvi, [
      ...
      { components, [
        ...
        { service_edge, [
          { url, "http://127.0.0.1:9011" },
          { exo_http_opts, [ { port, 9011 } ] }
        ]},
        { service_discovery, [
          { url, "http://127.0.0.1:9012" },
          { exo_http_opts, [ { port, 9012 } ] }
        ]},
        { schedule, [
          { url, "http://127.0.0.1:9013" },
          { exo_http_opts, [ { port, 9013 } ] }
        ]},
        { authorize, [
          { url, "http://127.0.0.1:9014" },
          { exo_http_opts, [ { port, 9014 } ] }
        ]},
        { protocol, [
          { url, "http://127.0.0.1:9015" },
          { exo_http_opts, [ { port, 9015 } ] }
        ]},
        { data_link, [

```

```
    { url, "http://127.0.0.1:9016" },
    { exo_http_opts, [ { port, 9016 } ] },
    { bert_rpc_server, [ {port, 9017 } ] }
  ]}
]}
}}
}}
```

Please note that IP addresses, not DNS names, should be used in all network addresses.

COMPILING THE RVI SOURCE CODE

Before a development release can be built, the source code needs to be compiled. Please see BUILDING.md for details on this process.

CREATING THE DEVELOPMENT RELEASE

Please note that a new release must be created each time the configuration file has been updated

Once a configuration file has been completed, a development release is created.

The difference between a development and a production release is that the development release needs the compiled files located in the source tree to operate, while a production release is completely self contained (including the erlang runtime system) in its own subdirectory.

Each release will have a name, which will also be the name of the newly created subdirectory containing the files necessary to start the release.

If a configuration file, `test.config` is to be used when building release `test_release`, the following command can be run from the build root:

```
./setup_rvi_node.sh test_rel test.config
```

Once executed (and no errors were found in test.config), a subdirectory called `test_rel` has been created. This directory contains the erlang configuration and boot files necessary to bring up the RVI node.

STARTING THE DEVELOPMENT RELEASE

The newly built development release is started using the `rvi_node.sh` tool.

In order to start the test release, named `test_rel`, created in the previous chapter, the following command is run from the build root:

```
./rvi_node.sh -n tes_rel
```

When a development release is started the erlang console prompt will be displayed at the end of the startup process, allowing for manual inspection of the running system.

Once the RVI node has been brought up, services can connect to its Service Edge and start routing traffic.

FAULT SEARCHING

TRAFFIC TARGETED FOR A SERVICE ON ANOTHER NODE IS NEVER FORWARDED

TBD. Check that static node's service prefix matches that of the destination service.

MORE