INFN-Cloud: Rules of participation

Executive summary

Since several decades, INFN runs and supports the largest research and academic distributed infrastructure in Italy, with a very large national data center (CNAF) and 9 other sizeable data centers, all interconnected at very high capacity through the GARR network. This infrastructure utilizes Grid or Cloud protocols to serve the needs of several tens of international collaborations in physics and in many other scientific domains.

INFN now offers to its users a comprehensive and integrated set of Cloud services through its dedicated **INFN-Cloud** infrastructure. The INFN-Cloud portfolio, available via an easy-to-use web interface, is defined upon clear users' requirements. It is based on composable, open-source solutions and can be easily extended either by the INFN-Cloud support team or directly by end users.

The INFN-Cloud infrastructure is based on a core backbone connecting the large data centers of CNAF and Bari, and on several federated sites connecting to the backbone. Services on the INFN-Cloud backbone are typically reserved for special purpose tasks (such as multi-site automated data replication), while the other INFN sites of the INFN-Cloud infrastructure can transparently run one or more of its services, thanks to flexible Cloud orchestration policies. Joining a site to the INFN-Cloud infrastructure is regulated by its Rules of Participation (described in this document) and must be approved by the INFN-Cloud project management board, composed by the INFN-Cloud Project Coordinator and by the INFN-Cloud WP Leaders. In case of special arrangements, the INFN-Cloud infrastructure can be transparently extended to other public or private Cloud providers to augment its capacity or solutions.

Access to the INFN-Cloud services is currently reserved to INFN personnel or personnel with whom INFN has established formal collaborations, such as research associates, although research agreements with other institutions are foreseen for the future.

Authentication and authorization for accessing all INFN Cloud services is enforced through the INDIGO-IAM federated solution, fully compliant with European Open Science Cloud (EOSC) and industry standards.

Document log

Issue	Date	Comment	Authors
0.1	2020 Feb 27	First draft	Massimo Sgaravatto
0.2	2020 Mar 10	Second draft	Marica Antonacci, Vincenzo Ciaschini, Giacinto Donvito, Doina Cristina Duma, Luciano Gaido, Massimo Sgaravatto, Vincenzo Spinoso, Sergio Traldi
0.3	2020 Apr 9	Third draft	Marica Antonacci, Vincenzo Ciaschini, Giacinto Donvito, Doina Cristina Duma, Nadina Foggetti, Luciano Gaido, Barbara Martelli, Massimo Sgaravatto, Vincenzo Spinoso, Stefano Stalio, Sergio Traldi, Riccardo Veraldi
0.4	2020 Sep 16	General refactoring of the document, trying to converge to a first version applicable only to INFN resource centers	INFN-Cloud WP4
0.5	2020 Sep 24	Integration of raised comments (after discussion happened in a dedicated meeting)	M. Sgaravatto
0.6	2021 Mar 16	• Integration of D. Salomoni and C. Duma comments. • Fix use of security team (WP4) wrt security incident team. • Specify that re-certification is needed after suspension. • Fix "policies and procedures" URL link • Fixed/removed some broken creferences • Added reference to "Scansioni di sicurezza e gestione degli incidenti su INFN Cloud" document • Fixed accounting references and endpoints	M. Sgaravatto
1.0	2021 Apr 13	First version (after internal review)	WP4
1.1	2021 Oct 04	Log retention: logs must be kept max 1 year	Massimo Sgaravatto
1.2	2022 Jan 19	Services on public and private networks Other minor changes	WP4

1.3	2022 May 16	Logging section revised (to address CCR requests) Other minor changes	WP4

Introduction

This document describes the requirements that must be met by a **resource provider (**also referred as **resource center** in this document) willing to participate in the INFN-Cloud federation, upon discussions with INFN-Cloud management. For the time being we focus only on INFN resource providers. Future releases of this “Rules of Participation” document may consider also the federation of resource centers belonging to other institutions (e.g. universities, other research organizations, etc.).

The technical participation of a resource provider in the INFN-Cloud federation is defined in terms of resources that it provides to one or more user groups and of service levels that it pledges to the INFN-Cloud federation. Resources that can be provided to the INFN-Cloud federation are typically computing (CPU and GPU) and storage resources.

The resources provided to the INFN-Cloud federation by a resource center are made available to the INFN-Cloud end-users who, using the “central” INFN-Cloud orchestration system, can instantiate different kinds of services on such resources.

Examples of services that can be deployed are:

- HTCondor batch clusters
- Kubernetes and Mesos clusters for container orchestration
- RStudio instances
- Virtual machines with storage (which can also be encrypted)
- Spark clusters
- Jupyter notebooks

Besides the resources integrated with INFN-Cloud, a resource provider can also provide other resources to local users or other user communities: the operations of the resources not integrated with INFN-Cloud are not bound to the rules, policies and procedures described in this document.

An explicit goal of the INFN-Cloud architecture and implementation is to simplify and minimize the impact of federating multiple service providers. INFN-Cloud supports the integration of many different Cloud stacks: a resource center is therefore free to choose the tools and frameworks used to manage its local resources, according to its own technology preferences.

Each resource provider in the INFN-Cloud federation keeps managing its services and resources. It has complete control on which user communities he wants to allow on its resources and which quotas or restrictions to assign to each user group.

A resource provider willing to join the INFN-Cloud federation will have to comply with all the established operational and security requirements, policies, processes and procedures[1]. All these policies aim at implementing service provisioning best practices and common procedures in order to guarantee the high quality of the INFN-Cloud provisioned services.

Information about the service level targets, the amount and configurations of resources provided by the resource center to the INFN-Cloud federation, the user groups that can use such resources, and any other information relevant with the resource provider participation in the INFN-Cloud federation are formalized in specific documents:

- The “Resource Center Operation Level Agreement (OLA)” document
- The “User Community Operation Level Agreement (OLA)” document(s)

The first document defines the minimum set of operational services and the respective quality parameters that the resource center is required to provide.

The “User Community OLA” document extends the “Resource Center OLA” document with information related to the provisioning of services and resources to a specific user community (there is one “User Community OLA” for each user community supported by the resource provider).

These are not legal contracts but, as agreements, they outline the scope and conditions of the collaboration to support the INFN-Cloud user communities.

Why joining the INFN-Cloud federation?

There are several advantages in joining INFN-Cloud, such as:

- to easily support and provide resources to many research communities, in a uniform and coherent way;
- to align the local policies and procedures with international good practices;
- to adopt best practices of multi-cloud federation for the benefit of the local users;
- to benefit from the services delivered through the INFN-Cloud central services, without the burden of installing and operating locally such services;
- to participate in possible e-Infrastructure projects as an INFN-Cloud compliant resource provider.

[1] https://www.cloud.infn.it/policies-procedures/

Compliance for resource access

The resource provider must enable the instantiation of services, requested by either administrator of services or end users, through the orchestration services of the INFN-Cloud PaaS. To do that, the resource provider must use one of the infrastructure management frameworks exposing interfaces compliant with the INFN-Cloud federated architecture.

The INFN-Cloud orchestration services can integrate resources managed by different Cloud stacks. At the time being, the integration of resource providers has been extensively tested using the following frameworks:

- OpenStack
- Mesos

We invite resource centers using other Cloud frameworks to collaborate with the INFN-Cloud developers and operators to have also such systems fully supported in order to increase the INFN-Cloud supported solutions.

Sites using OpenStack need to operate at least the following services:

- Keystone (authentication and authorization)
- Glance (images)
- Nova (computing)
- Cinder (block storage)
- Neutron (networking)

There are no other strict requirements on the infrastructure management frameworks (e.g. it is not mandatory to use a specific version of a system) but the list of supported frameworks, and their requirements (e.g. services that need to be configured, minimum versions, etc.) may change over time if needed (e.g. for operational or security reasons). Changes impacting the participation of resource providers which are already members of the INFN-Cloud federation (e.g. if the frameworks used to manage their resources are no longer supported by INFN-Cloud) will be managed on a case-by-case basis.

The instantiation of services using the core services of the INFN-Cloud PaaS is the only mechanism that must be supported by the resource provider, unless other specific agreements exist. Therefore the resource provider is not bound to provide INFN-Cloud users also with direct access to its resources at the IaaS level. As an example, for a provider using OpenStack as infrastructure management framework, it is not mandatory to provide access to its resources to the INFN-Cloud users through the OpenStack dashboard or through the OpenStack APIs.

Authentication and Authorization

INDIGO-IAM (Identity and Access Management)[2] is the service used in the INFN-Cloud federation to manage authentication and authorization. It is compliant with OpenID Connect/OAuth standards, thus allowing its ease integration with several off-the-shelf components, including many Cloud middleware frameworks.

In the INFN-Cloud IAM instance, users are organized in user groups, where each user group represents a coherent user community (e.g. an experiment or another scientific collaboration). A IAM user group can be composed of sub-groups, if a finer granularity is needed.

A resource provider joining the INFN-Cloud federation must first of all enable the INFN-Cloud IAM service as one of its supported identity providers. This can coexist with any other already existing framework used to manage local identities: the INFN-Cloud IAM service will manage the federated INFN-Cloud users only.

The second step is authorizing the relevant INFN-Cloud IAM user groups (the INFN-Cloud user communities that must be enabled by the resource provider will be defined in an agreement between INFN-Cloud management and the resource provider). If needed, the resource provider can also ban the access of specific users within an authorized user community.

If the Cloud middleware framework used by the resource provider supports multitenancy, different IAM user groups must be mapped to different local tenants. Considering as example OpenStack, this means that different INFN-Cloud IAM user groups must be mapped to different OpenStack local projects.

An adequate degree of isolation must be guaranteed between local tenants, in terms of security, performance and robustness. By default, unless there are some specific requirements, proper isolation must be provided also between the users of the same project. Therefore, for example, by default a user should not be allowed to destroy an instance created by another user of the same tenant.

The resource provider must also enable the 'ops' INFN-Cloud IAM user group, which is used for monitoring and testing purposes.

Resource allocation

Resources should be allocated by the Resource Provider considering as granularity the IAM user groups (which map to local projects). Considering again OpenStack as an example, this means that each

[2] https://github.com/indigo-iam/iam

OpenStack local project must be granted a quota of resources. In general, unless there are specific requirements, it is not needed to define quota at the user level.

The quota of resources that must be assigned to such user groups, is agreed between INFN-Cloud management and the resource provider.

The list of INFN-Cloud IAM user groups supported by the resource provider and the quota of resources assigned to such groups must be detailed in the "User community OLA" document(s).

Resource configuration

Different resource providers can configure their resources in different ways and can possibly apply different policies.

Examples of configurations that can vary among different resource providers include:

- the performance of the computing and storage resources;
- the level of reliability of the provided services and resources (e.g. redundancy level for ephemeral and permanent storage, that can be implemented via RAID, replication, etc.) and more in general of the resource provider infrastructure;
- the used technologies and the resulting QoS of the provided services (e.g. Infiniband vs ethernet for networking, magnetic vs SSD disks for storage, etc.);
- the CPU and RAM overcommitment factors;
- the maximum allowed size (in terms of cores, memory, disk size) of an instance.

When relevant, these configurations should be negotiated between the resource provider and the supported User Communities and then explicitly detailed in the "User Community OLA" document(s).

These different conditions can also be possibly used by the INFN-Cloud PaaS orchestration services in order to match the best federated resources for the deployment requests.

Supported end users' services

The high-level services that can be instantiated by the end users through the PaaS services on the INFN-Cloud infrastructure can be grouped into two main categories:

- "High level supported services", i.e. services instantiated using TOSCA templates implemented and supported by the INFN-Cloud developers and administrators.

- "High level customized services": all the other services, instantiated e.g. using custom TOSCA templates, defined and supported by users.

Another classification of end users' services is between:

- Services instantiated on a public network
- Services instantiated on a private network.

The former are services exposed on a public network, which are therefore directly accessible by the end users.

The services instantiated on a private network are instead services configured by the INFN Cloud central services through a "host proxy", and are accessed by the end users through a VPN.

The resource provider can:

- support only services instantiated on a public network
- support only services instantiated on a private network
- support both types of services (this is the recommended option, in order to be able to support all cases).

To support the instantiation of services on a public network the resource provider from

a technical point of view needs

- to enable the images and the flavors needed for the instantiation of the supported services;
- to provide the public IP addresses needed to access the service: usually it is necessary to provide one public IP address for each instantiated service;
- to enable the access to the needed network ports for such IP address(es): details are provided below in this document.

To support the instantiation of services on a private network the resource provider from a technical point of view needs:

- to enable the images and the flavors needed for the instantiation of the supported services;
- to configure the host proxies (usually one per tenant) to enable the configuration of the services by the INFN Cloud central services
- to configure the VPN servers (usually one per tenant) to enable the users to access the services

Networking

Due to the distributed nature of the INFN-Cloud federation, it is particularly important that each resource provider guarantees appropriate network capacity and connectivity.

The resource provider must have a reliable and performant network connection that has to be evaluated by INFN-Cloud Management in relation with the use cases and amount of resources that the provider is willing to support.

The endpoints exposing the services of the relevant infrastructure management framework (e.g. the ones needed to create new instances) must be reachable at least by the core services of the INFN-Cloud PaaS and by the INFN-Cloud monitoring servers.

Regarding the services instantiated by the end-users through the INFN-Cloud PaaS services, the ones instantiated on a public network are exposed through public IP addresses that must be provided by the resource center. The number of the public IP addresses provided by the resource center to a given user community must be specified in the relevant "User Community OLA" document.

To support the deployment of the high-level services on a public network and to allow their access by the users, the resource provider firewall has to be configured in a way to allow inbound access only to the following ports belonging to the public IP addresses allocated for such services:

- port 22
- port 80
- port 443
- any port greater than 1024 except for the following ones, which must be kept closed:
    - 1080 (socks proxy)
    - 1191 (gpfs) (udp+tcp)
    - 2049, 4045, 4046, 4049, 20048, 20049 (nfs) (udp+tcp)
    - 3260 (iscsi)
    - 3389 (rdp)
    - 5900 (vnc)
    - 5800 (jvr)
    - 10000 (webmin)
    - 6000 to 6023 (X11)

Any other port must be kept closed, unless properly documented and authorized. If a resource provider for some reasons needs to have other ports open, it must contact the INFN-Cloud security team (WP4) providing all the details related with the request. This must be done using the INFN-Cloud service desk (https://servicedesk.cloud.infn.it/) and then selecting "Technical support". All the motivations and details related with the request must be specified in the ticket. The INFN-Cloud security team must then take charge of the ticket without undue delays and will start a discussion with the proponent to determine the best solution.

It must be stressed here that both resource providers and INFN-Cloud require that the administrators of these services, identified as the users who created them, properly secure them by adopting the proper mechanisms (such as iptables, firewalls, OpenStack security groups, etc.) aimed at granting access only to the network ports that are strictly needed.

As explained in the "Scansioni di sicurezza e gestione degli incidenti su INFN Cloud" document[3], the public IP addresses provided by the resource center (needed for the services instantiated by the INFN-CLOUD users) will be monitored by the INFN-Cloud security incident team, in order to promptly detect possible vulnerabilities.

Support

The resource provider must address incidents involving resources and services hosted in its infrastructure and must satisfy all other service requests (change requests, information requests, etc.) issued by the INFN CSIRT and INFN-Cloud operation and security teams.

The expected response time, which depends on the impact of the issue, is specified in the "Resource Center OLA" document.

Service level targets

The minimum thresholds required for joining the INFN-Cloud federation for what regards:

- the monthly availability (defined as the ability of a service or service component to fulfil its intended function at a specific time or over a calendar month)
- the monthly reliability (defined as the ability of a service or service component to fulfil its intended function at a specific time or over a calendar month, excluding scheduled maintenance periods)

are specified in the "Resource Center OLA" document.

The availability and reliability of resource centers federated with INFN-Cloud are measured by the INFN-Cloud operation team using proper central monitoring systems.

A resource provider must enable this operational service monitoring. This simply means enabling the 'ops' IAM user group and allowing the INFN-Cloud monitoring servers to access the endpoints exposing the services of the relevant infrastructure management framework (e.g. to create test instances).

[3] https://www.cloud.infn.it/policies-procedures/

Security

Participating in the INFN-Cloud federation and allowing related inbound and outbound network traffic clearly increase the IT security risk for the involved resource centers.

The resource provider is responsible for taking appropriate measures to mitigate this risk, in collaboration with the INFN-Cloud security team. For this purpose:

- The resource provider must follow IT security best practices that include pro-actively applying software patches, updates or configuration changes related to security.
- The resource provider must comply with all INFN-Cloud policies and procedures[4]
- The resource provider must collaborate with the INFN-Cloud security team in monitoring the resources provided to the INFN-Cloud federation.

Obligations and responsibilities in terms of security of the different INFN-Cloud actors (service and resource providers, end users, resource administrators) are defined in the policy documents available at https://www.cloud.infn.it/policies-and-procedures.

Management of security incidents

The resource provider shall promptly report to the INFN Cloud Security Incident Team about suspected or confirmed security incidents that have known or potential impact or relationship to the resources, services, or identities within the INFN-Cloud federation. The detailed procedure that the resource provider must follow in case of a security incident is reported in the "INFN-Cloud policies and procedure" web page [5].

Traceability and logging

The infrastructure middleware and any other ancillary service deployed in the resource center must be configured to be able to identify the source of all actions (e.g. instantiation of a new virtual machine, deletion of a resource, etc.) and the individuals who initiated them.

End user instances must be configured to send their logs to an external server.

In case of a security incident, it must be possible to identify the instance where the problem happened, even if the instance was deleted in the meantime.

This logging information must be retained for a minimum of 6 months and for a maximum of 1 year. The INFN-Cloud security team may define longer periods of retention for specific services and/or operational requirements.

[4] https://www.cloud.infn.it/policies-procedures/

[5] https://www.cloud.infn.it/policies-procedures/

The resource provider shall use logged information only for administrative, operational, accounting, monitoring and security purposes. The resource provider shall apply due diligence in maintaining the confidentiality of logged information.

Accounting

The resource provider should collect and made available accounting information, to find out how the resources federated in the INFN-Cloud have been used and by whom.

The accounting system used in the INFN-Cloud federation is compatible with the system used in the EGI FedCloud.

This means that each resource provider must collect information about local resource usage (producing usage records in the format specified at: https://egi-federated-cloud.readthedocs.io/en/latest/federation.html#cloud-usage-record).

The usage records should then be sent, using the STOMP protocol, to the central INFN-Cloud accounting system (the endpoint to be used is: *accounting.cloud.infn.it:61615).*

The resource provider can use any tool to produce the usage records and to send them to the central accounting server. Resource providers using OpenStack can use the comprehensive solution (based on the CASO, collectd and APEL-SSM tools), documented at https://confluence.infn.it/x/xQLdAg.

Please note that only the usage records concerning the federated resources must be sent to the central INFN-Cloud accounting system.

Certification

Before joining the INFN-Cloud federation, a resource provider will go through a formal certification process, to verify its compliance with all the technical requirements. Details of such certification procedure are provided in the relevant document available in the “INFN-Cloud policies and procedure” web page [6].

If a site is suspended (e.g. because it failed to comply with one or more policies) it will have to go through the certification process again before re-joining the INFN-Cloud federation.

Withdrawals

[6] https://www.cloud.infn.it/policies-procedures/

A resource provider wishing to withdraw from the INFN-Cloud federation shall send notification to INFN-Cloud management in advance (at least 3 months); anyway before closing access to any of the resources that it has been providing, it must cooperate actively with INFN-Cloud management to investigate possible alternative arrangements.

A withdrawing provider must support the INFN-Cloud operation team in the migration of services, data, etc. to other resources. After such migration, the withdrawing provider must securely delete within 1 month all data (including backups) pertaining to INFN-Cloud and its users, except the ones that need to be kept as provided by law.

Violations

INFN-Cloud policies and procedures are designed to apply uniformly to all participants.

If a policy violation occurs due to unexpected or exceptional events, INFN-Cloud management must be promptly notified.

Resource providers who fail to comply with these policies in respect of a service they are operating may be removed (suspended) from the INFN-Cloud federation until compliance has been satisfactorily demonstrated again.

More details are provided in the “Resource Centre OLA” document.