



Adaptive edge/cloud compute and network continuum over a heterogeneous sparse edge infrastructure to support nextgen applications

Deliverable D7.5
Technoeconomic analysis
(I)



DOCUMENT INFORMATION

PROJECT	
PROJECT ACRONYM	ACCORDION
PROJECT FULL NAME	Adaptive edge/cloud compute and network continuum over a heterogeneous sparse edge infrastructure to support nextgen applications
STARTING DATE	01/01/2020 (36 months)
ENDING DATE	31/12/2022
PROJECT WEBSITE	http://www.accordion-project.eu/
TOPIC	ICT-15-2019-2020 Cloud Computing
GRANT AGREEMENT N.	871793
COORDINATOR	CNR
DELIVERABLE INFORMATION	
WORKPACKAGE N. TITLE	WP 7 - Impact-maximization activities
WORKPACKAGE LEADER	AEGIS
DELIVERABLE N. TITLE	D7.5 - Technoeconomic analysis (I)
EDITOR	Marco Di Girolamo (HPE)
CONTRIBUTOR(S)	Lorenzo Blasi (HPE), Patrizio Dazzi(CNR)
REVIEWER	Massimo Coppola (CNR)
CONTRACTUAL DELIVERY DATE	12/2020
ACTUAL DELIVERY DATE	30/12/2020
VERSION	1.0
TYPE	Report
DISSEMINATION LEVEL	Public
TOTAL N. PAGES	44
KEYWORDS	Federation, business model, governance, coordination, admittance

EXECUTIVE SUMMARY

This document results from the work of ACCORDION Task 7.5 and analyses federation models from previous research and standards. The goal is to apply some of the characteristics of the studied models to the ACCORDION federation of Edge resources.

The federation concept in ACCORDION has both technical and business implications. At the business level, a federation governance model should be designed that defines how the various actors interact and how the main federation processes work. At the technical level, all the components needed to support the defined governance processes should be identified. This document is mostly oriented at analysing the business aspects of federations, but also provides some technical insights that can be useful to design the ACCORDION system.

Business aspects concerning the governance model are first discussed in general terms (sect. 3) and then in more detail by analysing two relevant federation types: the model proposed by the FP7 EU project 5GEx (sect. 4.1), thought for a 5G network ecosystem, and the various models for cloud federations (sect. 4.2) emerging from both standard committees and (mostly EU) research projects.

Based on the previous analysis, the document then proposes some base ideas for the ACCORDION federation model (sect. 5). The proposed model, mostly based on the work from 5GEx EU project, designs a closed community with a good degree of trust between the members.

In the 5GEx federation each provider can act as an entry point for the whole system and for each service request it plays the role of service aggregator, providing both service composition and end-to-end SLA. Cooperative service provisioning may take the form of Task Forwarding, where a provider can delegate part of a request to other providers, or Capacity Sharing, where part of the resources from a provider can be shared to other peers in the form of "resource slice". For revenue sharing, 5GEx analysed several pricing schemes, but did not select a specific one.

From a technical standpoint, our analysis (sect. 6) is based on the latest standard from NIST, which provides a Cloud Federation Reference Architecture. Several federation processes are analysed, from membership, resource discovery and access, to monitoring and incident response. For each governance process, one or more architecture components are identified, such as Federation Entrypoint, Resource Repository, Federated Identity, Policy Enforcement Points, Monitoring, Event Recording, Logging, Filtering, and Aggregators.

DISCLAIMER

ACCORDION (871793) is a H2020 ICT project funded by the European Commission.

ACCORDION establishes an opportunistic approach in bringing together edge resource/infrastructures (public clouds, on-premise infrastructures, telco resources, even end-devices) in pools defined in terms of latency, that can support NextGen application requirements. To mitigate the expectation that these pools will be “sparse”, providing low availability guarantees, ACCORDION will intelligently orchestrate the compute & network continuum formed between edge and public clouds, using the latter as a capacitor. Deployment decisions will be taken also based on privacy, security, cost, time and resource type criteria.

This document contains information on ACCORDION core activities. Any reference to content in this document should clearly indicate the authors, source, organisation and publication date.

The document has been produced with the funding of the European Commission. The content of this publication is the sole responsibility of the ACCORDION Consortium and its experts, and it cannot be considered to reflect the views of the European Commission. The authors of this document have taken any available measure in order for its content to be accurate, consistent and lawful. However, neither the project consortium as a whole nor the individual partners that implicitly or explicitly participated the creation and publication of this document hold any sort of responsibility that might occur as a result of using its content.

The European Union (EU) was established in accordance with the Treaty on the European Union (Maastricht). There are currently 27 members states of the European Union. It is based on the European Communities and the member states’ cooperation in the fields of Common Foreign and Security Policy and Justice and Home Affairs. The five main institutions of the European Union are the European Parliament, the Council of Ministers, the European Commission, the Court of Justice, and the Court of Auditors (<http://europa.eu.int/>).

Copyright © The ACCORDION Consortium 2020. See <https://www.accordion-project.eu/> for details on the copyright holders.

You are permitted to copy and distribute verbatim copies of this document containing this copyright notice, but modifying this document is not allowed. You are permitted to copy this document in whole or in part into other documents if you attach the following reference to the copied elements: “Copyright © ACCORDION Consortium 2020.”

The information contained in this document represents the views of the ACCORDION Consortium as of the date they are published. The ACCORDION Consortium does not guarantee that any information contained herein is error-free, or up to date. THE ACCORDION CONSORTIUM MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, BY PUBLISHING THIS DOCUMENT.

REVISION HISTORY LOG

VERSION No.	DATE	AUTHOR(S)	SUMMARY OF CHANGES
V0.1	14/10/2020	Marco Di Girolamo	Proposed ToC
V0.2	29/11/2020	Marco Di Girolamo	First draft
V1.0	30/12/2020	Marco Di Girolamo, Lorenzo Blasi, Patrizio Dazzi	Final Version

GLOSSARY

BSS	Business Support System
CFRA	Cloud Federation Reference Architecture
DoA	Description of Action
EU	European Union
EC	European Commission
H2020	Horizon 2020 EU Framework Programme for Research and Innovation
ICT	Information and Communications Technology
IoT	Internet of Things
KPI	Key Performance Indicator
M2M	Machine to Machine
NIST	National Institute of Standards and Technology
NFV	Network Function Virtualization
OSS	Operations Support System
QoE	Quality of Experience
QoS	Quality of Service
SLA	Service Level Agreement
RBAC	Role Based Access Control
UI	User Interface
VIM	Virtual Infrastructure Manager

TABLE OF CONTENTS

EXECUTIVE SUMMARY	3
GLOSSARY	6
TABLE OF CONTENTS	7
TABLE OF FIGURES	8
1 Relevance to ACCORDION	9
1.1 Purpose of this document	9
1.2 Relevance to project objectives	9
1.3 Relation to other workpackages	9
1.4 Structure of the document.....	10
2 Federation concept in ACCORDION	11
2.1 Cloud federation perspective	11
2.2 Federation in ACCORDION	13
2.3 Business roles and actors	14
3 Aspects to evaluate	15
3.1 Federation governance	15
3.2 Admittance policies	17
3.3 Coordination model	17
4 Reference models	19
4.1 Model 1 – 5GEx	19
4.1.1 <i>Aggregation roles</i>	20
4.1.2 <i>Coordination models</i>	21
4.1.3 <i>Open vs closed community</i>	25
4.1.4 <i>Revenue sharing and charging</i>	26
4.1.5 <i>Technical impacts</i>	27
4.2 Model 2 – Cloud federations.....	27
4.2.6 <i>EU-funded Research projects</i>	28
4.2.7 <i>NIST Cloud Federation Reference Architecture</i>	29
5 ACCORDION federated business model	30
6 Inputs to technical design	39
6.1 Federation instantiation.....	39
6.2 Federation membership.....	39
6.3 Resource discovery.....	40
6.4 Resource access	40

6.5 Monitoring, Reporting, Accounting, Auditing 41

6.6 Incident response 41

6.7 Termination 42

6.8 Summary 42

7 Conclusions 43

8 References 44

TABLE OF FIGURES

Figure 1 - Baseline graphical representation to show the different 5GEx coordination options (Source: 5GEX deliverable D2.3, "5GEx Business and Economic Layer") 22

Figure 2 - Fully Centralized Push coordination model..... 22

Figure 3 - Fully Centralized Pull coordination model 23

Figure 4 - Distributed coordination model, push option..... 23

Figure 5 - Per-provider Centralize model, push mode 24

Figure 6 - Main functionalities for federation onboarding..... 35

Figure 7 - ACCORDION Resource State Diagram 40

Figure 8 - ACCORDION Federation Components Diagram 42

1 Relevance to ACCORDION

1.1 Purpose of this document

The present document is the result of the collaborative effort of all ACCORDION partners participating to Task 7.5. ACCORDION Description of Action (DoA) says: *ACCORDION aims at enlarging the edge resource & infrastructure pool by involving resource and infrastructure providers, such as telco operators or private companies, in some form of federation community. This task will study possible federation models by analyzing at least the following most important characteristics..[...].*

The document offers a review of the possible models of federation that can satisfy, for the specific context of ACCORDION, the main business criteria enabling the creation and tenure of a resource provider federation. This task investigates strictly business-related factors. Nevertheless, it also has an important connection with the technical design work carried on in WP2 (Task 2.3). The business model indications coming from Task 7.5 have indeed been shared with Task 2.3, and they hinted the specification of certain components in the ACCORDION reference architecture.

The work of Task 7.5 does not complete its cycle with the present deliverable. This report will have an updated version at month 24 (D7.9), that will refine the business model elements for the ACCORDION federation, on one hand by elaborating the external connections with potentials stakeholders, on the other hand by evaluating the possible feedbacks from the first pilot cycle that could provide technical recommendations impacting the business model viability. Task 7.5 will monitor the reference market, to ensure a proper response in case the risk Ext2 (see D1.2, Table 1) occurs.

1.2 Relevance to project objectives

The research topics analyzed in this document are quite important to achieve the ACCORDION project objectives. As far as the project Objectives (Ox) in the DoA are concerned, the Task 7.5 and the present deliverable directly tackle the objective KPI-O5-4, part of the impact maximization objectives.

1.3 Relation to other workpackages

WP7 is a standard horizontal workpackage, whose nature is peculiar and different from all the WPs directly tied to the design, implementation, and evaluation project cycle. Task 7.5 specifically has a connection to WP2 (Task 2.3), that it provides with indications in terms of certain functional components to be incorporated in the ACCORDION architectural design. It also has a loose inward connection towards WP6, since the outputs of the pilot cycle could suggest changes and/or refinements to the business model produced by Task 7.5.

1.4 Structure of the document

The present document can be seen as split in two main parts. The first part introduces the topic, describing how the federation concept is intended and declined in ACCORDION (section 2), what the key aspects to analyze for a resource federation in ACCORDION are (section 3), and the possible existing federation models coming from similar contexts that can inspire the definition of ACCORDION federation (section 4). The second part of the deliverable presents the actual choices done for this first version of the ACCORDION federation model (section 5), and what inputs the model could inject into the architectural design (section 6).

2 Federation concept in ACCORDION

2.1 Cloud federation perspective

Nowadays, Cloud computing paradigm is spreading in a very pervasive way. Cloud computing is everywhere. Infrastructure providers give their customers access to storage and computational resources that are paid for on an hourly or daily basis. Most application providers in their today's offerings almost always give the "cloud option" for (a subset of) their software. In fact, Clouds allow users to access computing services and resources on demand without having to buy their own infrastructures, and to pay only for what they use [1]. That is an advanced match of the Utility Computing paradigm.

Many cloud companies (mostly based in the USA) have developed their own platforms providing proprietary interfaces. Relying on such custom solutions actually affects interoperability, thus interconnecting with different platforms makes things difficult for customers who need the combined services or resources of multiple providers. This often results in users being locked into specific providers and platforms [2].

This issue has led to the introduction of interconnected clouds, also known as interclouds. Interclouds overcome single-provider approaches' limitations, for instance the lack of interoperability between platforms, peak customer demand, service interruptions, and quality-of-experience degradation.

An intercloud it's a large cloud comprising many smaller clouds, each having its own characteristics and serving different needs. An intercloud implementation could be any one or combination of [3]:

- hybrid clouds: private clouds access the resources of public clouds upon necessity;
- multclouds: applications use resources from multiple clouds, without any of them being aware of their participation;
- cloud federations, an interconnected set of heterogeneous public and/or private clouds from voluntarily participating users and providers.

As a matter of fact, the most successful type of Intercloud is cloud federations. Federations enable power-efficient, cost-effective, dynamic sharing of cloud resources and services. Federations can be based on specific Service-Level Agreements (SLAs) to ensure QoS and availability.

From a commercial perspective, Cloud Federations can act as a model for cloud providers that own multiple installations (for instance, in diverse geographical regions), acting in a way similar to Amazon's availability zones, with the purpose of balancing workload among them or for providing contextualized services. Even more importantly, it can enable groups of providers to offer communal services to final users, including or not a resource broker.

To realize a fully Federated Cloud vision, other aspects need to be developed, such as:

- collective resource provisioning and billing;
- cloud-across privacy, security and identity management;
- reliable mechanisms for data management;

- efficient and effective exploitation of diversity and heterogeneity of resources at all levels (compute, data and network);
- as well as, adoption of existing Cloud standards.

It is of paramount importance to ensure the full development of such capabilities to raise new business opportunities both for existing and new cloud stakeholders. This is especially true and important for the European market, typically characterized by players who are smaller than their American counterparts, that in this way could compete and gain shares in the Cloud market. A sustainable Federated Cloud market, with the related tools and technologies, may also allow enterprises and the public sector to rely on Clouds in a very flexible way and achieve the right balance among services provisioned in house and those consumed from external Cloud providers.

From a providers' perspective, the Cloud Federation model has the potentiality to overcome the traditional disjunctive between public and private cloud, hence, to include the necessary services across the entire computational continuum, realizing complex resource environments including large data storage, complex data centers, edge servers and IoTs. In fact, if on the one hand Federated Cloud markets are today being a reality, the emergence of Edge computing is actually expanding the scope of the "federated" concept.

Edge computing is rapidly growing and represents the major trend in the distributed computing area. Edge computing, is capitalizing the attention of both research and commercial communities, and is in fact at the core of ACCORDION project, being a key technology to enable the computing continuum that we are realizing in our project, actually allowing to rely on Cloud technologies even for nextgen applications.

Edge computing represents the first step towards the decentralization of Cloud computing and, as such, is bringing the concept of Federated Cloud one step forward. By means of the development of Edge computing, the concept of Cloud Federation is transformed in order to encompass an even more distributed approach which includes a diversity of cloud and edge offerings that lead to better performance which enable a wider diversity of application and services, complementarity to traditional cloud X-as-a-service models. Such a need for the exploitation of diversity is taken into due account in ACCORDION project, and finds its key enablement in many of the services realizing the ACCORDION platform.

Clearly, the development of a computing continuum technology compatible with the management of hardware heterogeneity requires to find new ways to optimally endeavor heterogeneous special purpose processing units without losing the advantages of abstraction of the utility-based models. These could be materialized by the development of:

- Novel virtualization tools,
- Heterogeneity-aware scheduling at resource and platform levels,
- Programming models permitting to seamlessly handle heterogeneity.

Moreover, taking advantage of hardware heterogeneity (together with specific developments for energy aware metering and scheduling) are key aspects to enable the advance towards energy aware, optimized and sustainable cloud continuum compute services.

2.2 Federation in ACCORDION

The federation concept in ACCORDION has a double meaning, according to the perspective from which we look at it (technical vs business-side). Although the present deliverable tackles the federation from a business perspective, it's useful to report here how ACCORDION defines a federation from a technical standpoint as well:

- At **business** level: association of resource providers willing to share their resources within a common ACCORDION technical framework
- At **technical** level: set of Miniclouds governed and controlled by a unique overarching ACCORDION orchestration entity

An ACCORDION federation coalesces together a set of technical resources, coming from both the edge and the centralized cloud layers, under the administrative control of different local infrastructure owners, choosing to dish out their resources to the federation. This federation must be governed by a set of rules and policies making it ultimately viable from a business standpoint to its participants, ensuring that joining the federation brings them a proper return for the availability of their resources, and the commitment to respect the agreed service level KPIs subsuming such sharing. The federation members should have a chance to take advantage from the porting of next generation applications into the federated domain. The active and committed attendance of infrastructure and resource owners to a federation requires to be properly incentivized, offering a scheme to grant a viable and effective sharing of federation revenues among its members. In other words, the main purpose for joining an ACCORDION federation is to make a profit out of the hiring of resources to support service delivery for ACCORDION customers. The project has the aim of investigating and finding out the most possible effective incentive mechanisms, to attract potential participants into its federations, increasing the share of providers looking at this as a crisp business opportunity.

A number of aspects define a federation governance model, specifying how members are admitted into the federation, what is the mechanism allowing to fairly share the revenues coming in thanks to the participation to the federation, how the federation is managed, how the access to federated resources is properly controlled, how the resources are exposed and made available to the whole federation. All these points will be better overhauled in section 3. A key point is granting the security of every interaction putting in touch an infrastructure owner with its federation peers. The federation as a whole requires that all the resources can be discovered and assigned to an orchestration process, thus making them usable to deploy client applications. For this to realistically happen, the federation mechanisms must ensure full isolation of the shared resources from any other sensitive section of the member's infrastructure, to safeguard the integrity and privacy of the resources not shared with the ACCORDION federation from any unauthorized and/or malicious access and manipulation. Clear rules and non-ambiguous, unmodifiable business agreements are a key enabler, as well as a consistent mechanism to manage revenue sharing and to resolve disputes.

2.3 Business roles and actors

A federation typically involves a certain set of actors and actor roles. Such roles identify actors playing as providers for certain elements of the federation, and at least one role identifying the federation customer, who is an external role, nonetheless fundamental to motivate the existence of the federation and to make sense of its business model.

In the ACCORDION federation, we can initially identify three internal provider roles, and one external customer role. We briefly summarize them hereafter.

- *Application vendor/provider.* This is the actor playing the external customer role for the ACCORDION federation. ACCORDION mission is delivering to an application vendor the resources, distributed across the edge to cloud continuum, necessary to optimally deploy and execute its application. The point of contact between application vendors and the ACCORDION federation is one of the aspects defined by the federation model, specifying the kind of entry point for the external customer to require and obtain the resources from the federation
- *Federation owner.* This actor can be identified as the founding member of an ACCORDION federation. In more practical terms, it is the entity which first creates and sets up a Minicloud in the ACCORDION federation. The federation can, at a general level, include both members running their own Minicloud (through a mini-VIM layer as described by the DoA), and members only offering resources but not hosting a Minicloud. However, at least one Minicloud must exist in the federation to allow the execution of the resource allocation and orchestration functions defining the ACCORDION platform. Hence, there will always be a federation owner.
- *Minicloud owner.* This actor represents an active member of the federation, which, after joining the federation, deploys and exercises its own Minicloud. The Minicloud can manage resources owned by the same actor as well as resources of other entities not running their own Minicloud
- *Resource provider.* This actor represents an active member of the federation which makes some of its resources available to and shared with its federated peers, but does not run its own Minicloud. Consequently, its resources are put in touch with the federation through one of its federated Minicloud owners.

This set of roles is a general representation of the actor ecosystem that can be expected to live and act around the ACCORDION federation. The distinction between the three roles internal at the federation (federation owner, Minicloud owner, resource provider) is nonetheless mostly visible at a technical implementation level, where the actual capabilities and functionalities offered by the different actors are not the same for all. For the purpose of the current document, i.e. the description of a general initial business model for the federation without a detailed vision of the underlying technicalities, the three federation member roles are not as different, or, in other words, the characteristics differencing them are not strongly relevant to evaluate the business model aspects that we are going to explore. So, in the following of the current report, we will mostly use the terms “resource provider”, “federation member” or “federation peer” to generically refer to entities part of the federation where the specificities of its role do not influence the options and choices available in evaluating the business model.

3 Aspects to evaluate

3.1 Federation governance

The federation can be considered as a community of providers who are normally business competitors, nevertheless they find a common interest and profitability in joining their resources according to a collaborative scheme. This model is by no means new, and it's particularly emerging in the recent years within the telecommunication market. Such a "coopetitive" model is proving to be an effective solution to reduce the overwhelming cost of infrastructure maintenance, particularly heavy in the current transition from 4G to 5G networks, that is critically undermining the operators' and providers' profitability margins.

A similar concept looks very promising when ported to the domain covered by ACCORDION. The shift from centralized clouds towards the edge brings at the spotlight a number of local resource providers, who however can be challenged to support the compelling requirements posed by next generation applications. Hence, even in this scenario a collaborative federation of resources may turn out to be a very promising solution. Clearly, sharing own resources with other entities who are at the same time potential competitors requires to meet a set of key pre-requirements to ensure that the federation does not raise issues of protection for the exposed resources and of conflicts related to their shared utilization for the delivery of services to common customers. Turning such issues around means having a clear, solid and efficient governance mechanism.

A governance model must ensure that a baseline of policies and rules is set out, prescribing how the federation handles any possible situation. The specification of policies and rules should be flexible, to adapt itself at the change of scenarios and business relationship. In particular, flexibility is needed if the federation wants to handle differentiated levels of trust among its members, a feature that can help to extend the federation coverage. An appropriate duly implemented governance model is a determinant factor to exploit the technical capabilities offered by the ACCORDION framework, to increase the effectiveness of the whole ecosystem, encourage the collaborative competition among the involved providers, and ultimately increase the dependability of the service offered to ACCORDION end users. A suitable governance model increases the federation business sustainability.

The governance model encompasses (1) the general configuration of business relationships among the federation participants (2) the punctual business agreements acting among the federation members, that every candidate participant must abide by as a mandatory condition to join the federation itself. The governance model is implemented through a set of policies and rules, encompassing, among the other items, the tenets underpinning the collaboration among federation members, the macro-level management of the federation, the business models, the procedures for admittance and withdrawal to/from the federation.

The collaboration tenets should allow to define, as precisely as possible, at what extent the federation members are requested to keep aligned both their technical and business related procedures. From a technical standpoint, the collaboration tenets should outline a clear set of technical pre-requirements to be mandatorily met by each of the federation members. Each candidate participant must accept to implement the needed ACCORDION capabilities, by deploying on its resources the fundamental ACCORDION software

components and implementing the required endpoints to make its resources accessible. The governance model does not prevent of course to extend the technical level of federation by increasing the amount of shared information, e.g. opening own resources up to deeper monitoring, making visible their utilization level, offering more reliable endpoints, and so forth. All such extensions can overall allow to ACCORDION a more effective allocation of resources to applications, and ultimately an improved service level. Apart from the technical requirements raised by similar agreements, there is an important business dimension ruling out what kind of information the federation members are willing to expose and share, and to what counterparts they are available to grant this information transparency.

As a minimum, the collaboration tenets must support the basic functions required to operate a federation, complying with the common ACCORDION resource description model, prescribing resource charging rules, agreeing a process to manage Service Level Agreements. On top of these essential functions, a federation may decide to implement extra business logic, for instance more sophisticated revenue sharing mechanisms, or penalty rules for breakings SLAs.

The governance model specifies the entity in charge of overall federation management and decision making. The model can be implemented through a centralized governance entity, or in distributed mode with responsibilities shared among all the federation members. In the centralized mode, the central administrator entity can be a third party trusted broker; this can happen where the market weight of the federation members is reasonably balanced, and there is not a de-facto evident leadership. This appears as a reasonable choice for the ACCORDION federation, that may be envisioned as a community of mostly small-medium providers. It is anyhow possible that one of the members (or a restricted subset of members, typically the biggest ones) is assigned a leading role to define the business rules and take control of the resource allocation process. The centralized mode is more common in closed federations (see sect. 3.2), whereas open federations are more prone towards a distributed governance.

A typical set of general governance rules may include:

- Compliance to code of conduct, business agreements, contractual templates and policies
- Compliance with ACCORDION technical procedures and definitions
- Commitment to share and expose a catalogue of available resources in the format required by ACCORDION
- Commitment to install and run the ACCORDION software on the shared infrastructure, and to configure the shared resources as dictated by the federation
- Compliance with non-disclosure rules outside of the federation for the sensitive data (customer data, pricing lists, system monitoring data)
- Commitment to enact in the member's infrastructure the ACCORDION monitoring facilities, allowing access to the performance measures to the federation managing entity, at the extent needed to properly ensure a common SLA enforcement

Closed federations, with higher trust level among members, may incorporate additional business rules, subsuming tighter commitments as well as the deployment of additional technical components. Some examples:

- Commitment to refrain from pricing discriminatory practices, and/or unfair trading
- Tighter commitments in terms of transparent information as far as resource availability, utilization and service assurance are concerned
- Full abiding by the federation business rules, also for customer management and gains/penalties sharing
- Commitment to respect the agreed guidelines in terms of pricing, charging and gain sharing
- Adherence to rules regulating admittance and withdrawal of members
- Adherence to the common coordination model

3.2 Admittance policies

A key aspect of federation business is the set of procedures ruling out how a resource provider can be admitted to become part of the federation itself, and, on the contrary, what rules and procedures must be followed for a member to leave the federation.

As a general rule, a federation can opt to be open or closed, depending upon the existence of pre-conditions to be met for a new member to enter the federation. An open federation might even simply not have any admittance procedure in place at business level (from a technical standpoint, there will anyway be some pre-conditions to be abided, at least the acceptance to install and execute the ACCORDION platform software on the entering provider’s infrastructure). In similar cases we talk about an “implicit” admittance procedure. The mobile telecommunication market is one of the best-known cases of open model federation.

In closed federation models, instead, given the increased level of information sharing agreed among the members, it’s impossible to admit a new member if a proper level of trust and cooperation cannot be ensured. Even more, being admitted to a closed federation imposes to be compliant with the stated collaboration tenets, including all their business as well as technical sides. Consequently, the federation has to define a clear procedure to ascertain and certify the new entrant’s conformity to the collaboration tenets of the federation. A common reference for a closed federation is the airline market, which also shows a degree of oligopoly governance, dictated by the biggest players in the market.

3.3 Coordination model

In the ACCORDION federation, context, as “coordination model” we intend the way through which a resource allocation request coming from outside ACCORDION is handled across the federation workflow, encompassing (in general) steps of resource discovery, resource negotiation and resource assignment. The workflow involves the identification of different business roles, and can be resolved according to different philosophies. In a generic coordination model, the differentiating factors are centralization or distribution of control, but also different flavors of centralization, and diversity of initiation modes (push vs pull). The choice of a coordination model has a significant dependence upon the type of federation being enacted (open or closed).

Coordination models normally encompass two phases: a “publishing” phase and an orchestration phase. The publishing phase covers the exposure and sharing among the different providers of information relevant to the potential resource delivery requests coming into ACCORDION. The orchestration phase is triggered by an actual resource delivery request, and uses the information collected during the previous phase.

A centralized coordination model can assign the central coordinator role to one of the federation peers, or to a third-party business entity. The central entity can have an aggregated view of the whole federation in terms of available resources, which can allow a more streamlined process of resource orchestration. In a distributed mode, the typical modality is a cascade process, where each provider starting from the first taking a given service request evaluates its ability to fulfill the request with its own resources, and, in case it’s not able to complete the request, forwards it to its neighbor provider in the federation.

The centralized coordination model can in turn be fulfilled by a unique central entity, or from multiple providers assigned to different types of service requests. In the ACCORDION case, this latter is not a very likely option, given the not so wide range of service types that the system is supposed to serve, probably not motivating such an articulated mesh of coordination.

In terms of initiation mode, this differentiation refers to the publishing phase. In a push mode, the federation providers share upfront not only their available resource catalogue, but also a set of actual SLA offerings, where given, even pre-packaged, resource sets are upfront associated to a range of KPI commitments and optionally even prices. In the pull mode, the provider exposes a set of available resources with generic QoS indicators, but the actual service offerings are created only in response to an actual service request.

The aforementioned features of coordination models can be combined in different ways, generating a number of possible options. A detailed analysis of such combinations of performance and applicability to different use cases, even in terms of model scalability, has been performed by the 5GEx project¹.

¹ Deliverable 2.3 - 5GEx Business and Economic Layer

4 Reference models

4.1 Model 1 – 5GEx

The 5GEx project defined a set of fundamental business roles for a 5G network ecosystem resembling at good extent some key features that may also be found in the ACCORDION federation. 5GEx defined the key attributes of what it calls a *5G community*, analyzing how the specific market conditions and the business structuring of the whole ecosystem might affect them. In the 5GEx community, the glue linking together the federated members is the co-opetitive service model, enabling different providers to cooperate while at the same time remaining competitors with each other. The final business goal of the federation is the delivery of 5G services whose provisioning spans different administrative (and technological) domains. To enable such a community for running actual business, clear collaboration tenets must be established, and each community member must fully adhere to them. Non ambiguous rules and policies can guarantee that the community business is sustainable, offer real incentives to be part of the federation by allowing to realize visible, concrete gains from this cooperation, and ensure that the co-opetition constantly goes on fair trade and competition pillars, with no member taking improper advantage from its stance in the federation. The collaboration tenets must rule out all the key aspects tied to level of cooperation, revenue sharing and liability management in case of service level breach.

The actual relationship among the federated members and their level of trustworthiness are clearly paramount elements influencing the community operation. This spawns a first important distinction among *open* and *closed* communities. In an open type community, the main effort of members is focused on the technical alignment of members. In a closed community, policies are tighter, and much focus is set on trust and rule compliance. Both the types (open and closed) can nonetheless exist under a same federation model, that can be implemented in both of these fashions.

The 5GEx federation defined a number of actor roles able to exhaustively cover the use cases envisioned in a co-opetitive 5G market scenario, including roles of NFV operation support and service aggregation essential to make such a 5G ecosystem adequately scalable. Additional elements considered by the 5GEx analysis are the power balancing among the federation members, the trustiness degree acting among the federation members, and the diversity of services generated and delivered by the federation.

5GEx created a taxonomy of business roles requested by a 5G federation, defining five main role categories, starting from the final customer outside the 5GEx ecosystem, down to the customer front-end provider, the general service provider (which is also an internal client in the 5GEx ecosystem), support, supplier and resource/service aggregation roles. The external customer in the 5GEX context is an enterprise, so the 5GEx business is a wholesale class. The 5GEx external customer is a reseller of network-based services to its own customers, ranging from basic connectivity access to high value-added network services.

In the 5GEx model there is not a broker attested in front of the whole 5GEx federation. To access the resources and services offered by the federation, the external customer goes through one of the federated providers acting as front-end, unique entry point to access the service catalogue and purchase the desired ones. The 5GEx federation does not impose to every of its members to act as customer front-end. A federated

provider can choose to just be a supplier for its federation peers, without taking charge of hosting and managing its copy of the federation service catalogue, and of looking after the direct service delivery to an end customer. This avoids to a federation member the burden to deploy and execute also the business layer components of the 5GEx software stack. More precisely, the front-end providers have to implement what in the 5GEx architecture is dubbed interface /1, an access point where the customer can find the description of available services, and offering the functionalities needed to negotiate and finalize the purchase of such services. This federation customer entry point has also the purpose of discovering and negotiating the best possible offer for the incoming customer, when there are multiple alternative offerings from the federation members, or when it's not possible to fully meet the service request submitted by the customer (both in terms of technical offer and pricing), and a best effort option must be made up and offered.

The internal 5GEx federation providers are granularly classified according to the type of resources/services they can offer. They range from infrastructure providers (the actual typology found in the ACCORDION federation) to pure connectivity providers, exchange point providers, network services providers, and application-level service providers. Pure suppliers in the 5GEX federation (not buying any resources by their federation peers) are located at the lowest hierarchy layers, inside the NFV layer, hence NFV infrastructure and Virtual Network Function suppliers, along with related deployment and operation management roles.

4.1.1 Aggregation roles

Aggregation roles are also essential to the 5GEx federation. There is a Service Aggregation role, in charge of wrapping an end-to-end service out of the partial services or resources provided by other federation providers. The 5GEx federation foresees different possible types of aggregator role:

- *Service Aggregator*. This role performs aggregation of resources/services from different federated suppliers (service providers), realizing a sellable end-to-end object from the individuals service delivery components delivered by the underlying suppliers. The value of the aggregator role is in this unifying function, for instance acting as aggregator of services/resources located in geographically dispersed places. It's fundamental for the service aggregator to ensure full business neutrality with respect to all the federation members: its decisions in performing the service aggregation must never be biased by improper own interests, nor by unbalanced, privileged business relationship with certain members of the federation. A messaging service provider is a suitable example.
- *Hub*. Compared to the Service Aggregator, the hub is not a third party to the federation. It is a federation member itself, acting as customer access point for certain service types, geographic areas, or based on other parameters splitting the federation entry points for the external customers. A transit or exchange service provider is a suitable example.
- *Federation Facilitator*. This role is conceived to facilitate particular business functions in line with the federation agreements, policies and governance rules. The items suitable to be handled by a facilitator are for instance cost and revenue sharing, information sharing, sometimes even the service allocation. This role can be covered by an actual entity who also fulfils the hub role.

4.1.2 Coordination models

Among the main novelties delivered by 5GEx there is an interesting study on possible coordination models for multi-provider service discovery, negotiation and fulfilment capturing all the involved resource domains. In a federation of this type, the coordination model is one of the key elements to characterize how the whole business is carried on. As “coordination model”, 5GEx means the way how a multi-party, end-to-end service is discovered, negotiated, built up and contracted within the federation. 5GEX tagged the possible coordination models as **fully-centralized, distributed** and **per- provider centralized**, to whom an additional variant is applied based on the push or pull kind of approach used in the federation.

First key distinction is the one between centralized and distributed coordination. In the centralized case, there is one central entity holistically taking over the role of service aggregator. Such central entity can in general be one of the federation members as well as a 3rd party aggregator, not owning resources and not acting as a resource provider, but running the needed parts of the 5GEx software framework, namely the multi-domain orchestration business functions and the so-called *I2-C* interface. The central aggregator has the goal of composing the requested service in a way allowing to enforce the needed Service Level Agreement, in a context where the service is not normally simply composable by wiring a sequence of neighbor cascading domains.

The centralized coordination can in turn be *Fully* or *per-Provider* centralized. In a Fully Centralized mode, the aggregator entity is unique. In a per-Provider Centralized mode, the role can be taken by more federated providers, depending upon the requested service type and the customer specifics. This latter mode is especially relevant in case you want to implement hierarchical nested (clusterized) models of federations. To enable such more complex hierarchy, there are local fully centralized aggregators in the individual clusters, whereas at the upper inter-cluster level they act as a full mesh.

In the distributed coordination fashion, each of the federation members is paired with another neighbor one, and, whenever the service composition flow goes through it, it communicates with its neighbor peer, ideally in a purely cascading mode.

Whatever mode is chosen, each coordination model actually runs through two distinct phases: a *publishing* phase and a *service composition* phase. The publishing phase specifies and dictates the depth and granularity of the information sharing among the federated providers about the acting service catalogue offering. The publishing phase is propaedeutic and predecessor to the execution of the service composition, which is triggered any time a new service request is issued by an external customer.

Another key distinction within a coordination model is between push and pull modes. This difference concerns the depth and quality of the information exposed by resource providers in a publishing phase. In a push model, federated providers are active in background preparing upfront (preliminary to incoming orders) service wrappings and exposing related Service Level offers. In the pull mode, instead, the information exposed during the publishing phase is more generic, and only qualitatively describing the available resources/services, the offered Quality of Service, possibly a generic price range. Real SLA is defined and offered only in response to an actual customer request.

The options described above can be combined to create a matrix of six overall possible generic coordination models, a kind of 3x2 matrix where the columnar dimension is given by the push vs pull mode option. The figure below is a baseline graphical representation to show the different options. The square vs circle shapes refer to different offered resource types, being 5GEx a heterogeneous resource scenario.

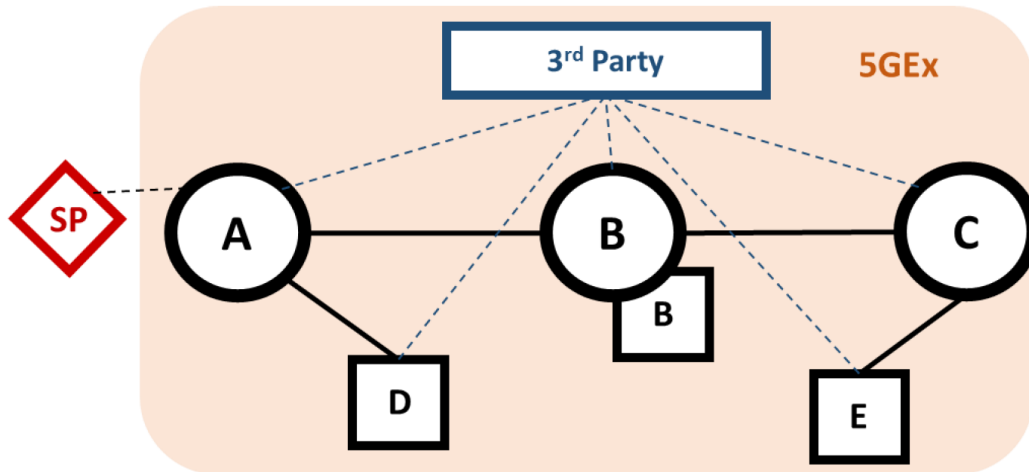


Figure 1 - Baseline graphical representation to show the different 5GEx coordination options (Source: 5GEX deliverable D2.3, "5GEx Business and Economic Layer")

The 3d Party is existing only in a Fully Centralized model, where it fulfills the aggregator role. SP is the external customer, assumed to have a settled legacy business relationship with one of the federation providers (A), which hence takes the role of primary access point for SP. A will be the primary access point for SP even in cases where the 3rd Party aggregator exists. Let's have a glance at how a coordination model is declined across the ecosystem actors. We take for instance the case of a Fully Centralized Push model.

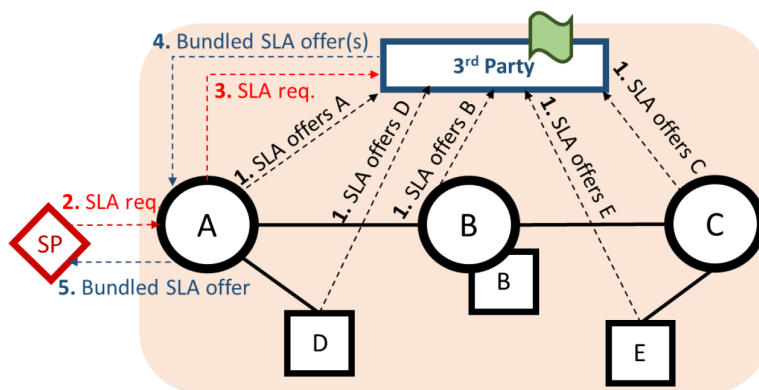


Figure 2 - Fully Centralized Push coordination model

4.1.2.1 FULLY CENTRALIZED, PUSH COORDINATION MODEL

During the publishing phase, all the providers (step 1 in Figure 2) expose to the aggregator their offered resources/services, properly formatted according to the syntax and semantic of 5GEx service catalogue. Such entries include the SLA specification, committed QoS indicators, price and time validity of the offer. Next

comes the service composition phase. Here the aggregator uses its technical tools (e.g., the abstract topology view) to match the offered resources/services with the end-to-end service request coming from SP (step 2-3). The aggregator normally comes up with more than one alternative composed offer, and returns this bundle to the customer SP (step 4-5)

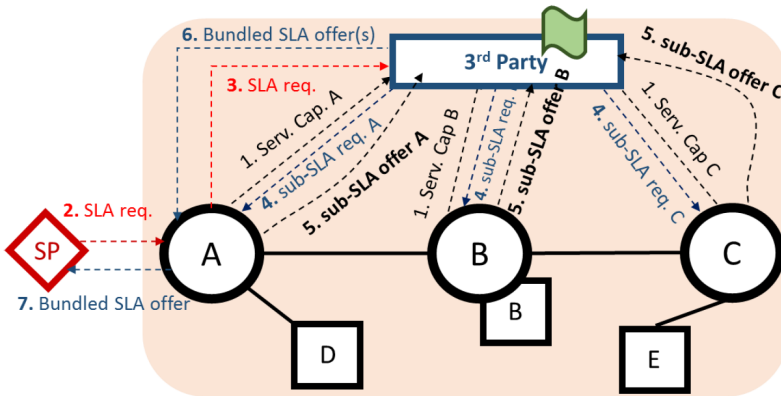


Figure 3 - Fully Centralized Pull coordination model

4.1.2.2 FULLY CENTRALIZED, PULL COORDINATION MODEL

In the case, represented in Figure 3, instead of a Fully Centralized Pull model, the aggregation step is not directly performed by the 3rd Party, since he does not have a full set of shared information sufficient to complete the operation in one step. Hence, as step 4 the aggregator dispatches sub-SLA requests to the providers candidate to participate the composed end-to-end service, each sub-SLA concerning a given part of the service (e.g., computational resources). The providers return their sub-SLA's in step 5, and then this case rejoins with e push one.




Figure 4 - Distributed coordination model, push option

4.1.2.3 DISTRIBUTED COORDINATION MODELS - PUSH OPTION

Distributed coordination, represented in Figure 4, is founded on *bilateral cascading* to propagate exposed resources/services and SLA offers through the federation. SLA offers are exchanged between direct neighbors during the publishing phase. The receiving provider can decide to bundle its neighbor's offer with its own, and to propagate this superset in turn to its next peer. Step 1-3 represent here the publishing phase. After

step 3, the federation member A has the information to build a service chain down to provider E. This upfront activity of pre-bundling enables A to respond immediately to a service request coming from SP.

In a pull mode option, additional messages will go among federation peers even during the composition phase, with the same sub-SLA mechanism described for a centralized model.

A distributed coordination model can suffer scalability issues when the number of federation peers is very high (flooding effect). This can be mitigated by setting proper rules taking advantage of the peers distribution, e.g. by setting up a maximum number of peers that can be crossed by the bundled offers forwarding. Clearly, a complex trade-off should be applied to avoid scalability is obtained at the expense of efficiency/effectiveness (too many options becoming unavailable for the service provisioning).

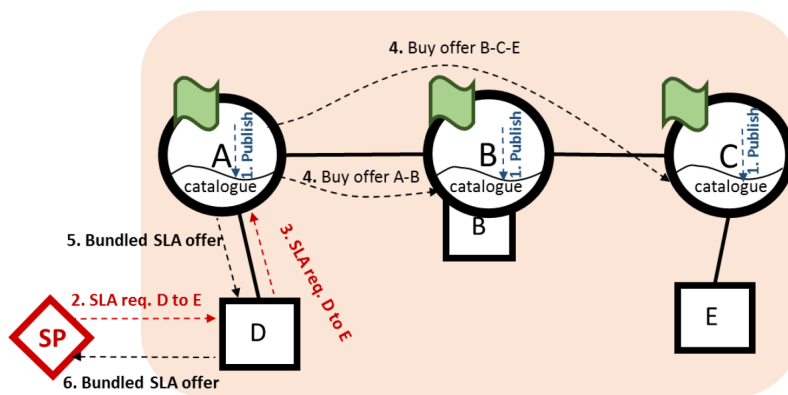


Figure 5 - Per-provider Centralize model, push mode

4.1.2.4 PER-PROVIDER CENTRALIZED MODEL (PUSH MODE)

In this final case, represented in Figure 5, the publishing phase occurs separately in each cluster (A-D, B, C-E). The local cluster aggregators (A, B, C) receive SLA offers both from inside their own cluster and from other clusters. The local cluster offers are stored in a repository (catalogue) accessible by every aggregator in the federation. The front-end provider, upon reception of the customer request forwards it to its neighbor cluster aggregator (A here), who takes over the role of composing the service, based upon all the entries stored in the different cluster repositories under its visibility. As usual, in the pull mode case there is an additional exchange of sub-SLA offers among the clusters during the composition phase.

The analysis performed within 5GEx concluded that the best coordination model was a Per-Provider centralized model combined with a distributed cascading model. The Fully Centralized model was discarded for business reasons strongly inherent to the telecommunication market. The largest operators, surveyed not only within the project itself, but also through external events, workshops and panels, expressed a general wide disagreement on the option to delegate the control of their service aggregation to a third party. They preferred to take the role of aggregators on behalf of other, more specialized providers, that is fairly well mapped into a Per-Provider centralized model. On the other hand, a pure Cascading Distributed model is simple, and can be designed to fit the actual business relationship hierarchies. However, it is less performant because it prescribes that each peer-to-peer transition along the service chain subsumes a negotiation step to respect the cascading path. This cascading flaw might be overcome by applying a full-

mesh coordination model (i.e., the extreme case of Per-Provider centralized model), that does not foresee a cascading flow, allowing each party to straightforwardly choose its negotiating entity in the chain. Nonetheless, the analysis showed that this solution sparks a significant business overhead, requiring to settle and maintain a high number of business relationships instead than just the ones with its own neighbors.

5GEx also performed an accurate scalability study on the different federation models. As a result, it highlighted that:

- Distributed models don't scale as well, being them impacted by a number of parameters, at least in an ecosystem like 5GEx particularly variegated in terms of service types and Quality of Service classes. Also, control messaging traffic may quickly grow (exponentially) with the addition of new parties.
- In terms of push vs pull, pull mode is convenient in the publishing phase, reducing the size of exchanged SLA offers. Push mode, in contrast, speeds up the service composition phase. Consequently, pull mode is good for young marketplaces of small size, whereas a push mode can better serve more mature and complex scenarios.
- The Per-Provider centralized model has the best scaling of all, even far better than a fully centralized model as long as the size and complexity of the federation increases. Its headstart comes from the reduced distance that messages cover during the publishing phase, being them bound to reach local cluster aggregators.

4.1.3 Open vs closed community

The distributed multiparty format is better fitting an open community type, where there is not a high trustworthiness degree among the federation members. This happens because business negotiations are substantially bilateral, not demanding the additional spreading of information nor business alignment. This approach is well suited to serve ephemeral supplier/provider business agreements. In this scenario, the community governance is fully distributed, and the admission policy is fundamentally demanding technical compliance alignment from the entrants. This kind of situation is, at first glance, the most similar to depict the ACCORDION ecosystem, at least in its first definition.

In contrast, fully centralized formats are better fitting a closed community, where each member is substantially trusted by its federated peers. In this scenario, typically applied to large markets with a growing business volume, widespread information sharing can enable a better utilization of resources and better support economies of scale. This must be supported by a deeper extent of information sharing. In such a scenario, the fully centralized model enables more sophisticated and flexible revenue sharing agreements and pricing schemes, compared to the essential bilateral charging and pricing schemes allowed by a distributed model. Here, the admittance policy is of course more complex, and scalability may become problematic with the skyrocketing of service types, and a potential differentiation of trust levels among members.

Partially centralized formats are best suited to populated ecosystems showing a non-flat business hierarchy, with a blend of peer-to-peer and client/provider relationships. Similar scenarios are realistic in the 5G enabled marketplace, for instance you might think about a "higher" layer closed community whose members

have open communities in the lower, resource closer layers. A community like this will likely hand over its governance to the upper tier members, and will be very well fitted by a per-Provider centralized model.

4.1.4 Revenue sharing and charging

The financial side of federations is easily enough shaped and handled in open federations. It gets more articulated in complex and more heterogeneous federations, where the governance is not equally distributed and the business weight of members is not balanced. 5GEx did not elaborate a “final” unique pricing and charging scheme, it rather provided some perspective into certain factors influencing the financial management of complex federations. For instance, if the federation is externally observed as a unique business entity (a kind of 5G virtual operator), on top of individual service request pricing there can be an offline clear-out mechanism, redistributing part of the revenues on a periodic baseline to fit the actual resource availability provided by the different members. In this hypothesis, revenues can be shared in the federation according to standardized mechanisms like the Shapley values or similar (5GEx made a detailed analytic elaboration on such cases). Alternatively, each provider can be the owner of business relationship with its own assigned customers, but revenues can as well be handled at federation level, recognizing a seamless reward to the provider “bearing” its customer as non-evident item in the service remuneration. This resembles the revenue sharing mechanisms applied in the airline sector to handle code-shared transportation services. The same principle can apply to penalties, which can be as well shared in case of service level breaches.

An ultimate outcome of the extensive analytical investigation performed by 5GEx is the following: in a federation fully driven by selfish objectives of reward maximization, the federation experiments a convergence towards a unique equilibrium state, where the profits granted to federation members substantially equal the case of joint business cooperating federations. In other words, there is an optimal solution somewhere satisfying both the cooperative and selfish federation cases, and similar provider subgroups within the same large federation. Other learning was that a federation based on capacity sharing can turn out more profitable than one based on task forwarding principle, but it has no way to align the optimization of financial profit with delivered quality of service, being it oriented to “give away” QoS for profit.

In terms of pricing schemes, the 5GEx federation study concluded that a “unary” price scheme was not really viable. Static pricing of a service can only expose a range, and the exact price of a service instance has to be evaluated and negotiated dynamically at provisioning time, since the actual instant availability of a given resource influences its economic value. 5GEx created anyhow a taxonomy of service categories and specific types, coupling each one to a specific charging model.

The whole work was based on a study on the relationship between network traffic flows and money flows. The pricing analysis also took into account the impact of costs stemming from the federation settlement and the need to maintain operational coordination between the federation members, these latter mainly capital expenses. Hence the total price comes from a combination of infrastructure operational expenses, service creation and operational management, costs of coordination with federated peers, and of course a needed profit margin. The conclusion was that 5GEx federated peers always have to perform a punctual evaluation

of cost/benefit ratio. In case of 5GEx, profits can have spot peaks due to particular situations where for instance the geographical location of a provider puts him in a headstart. Other considerations are tightly coupled to the network nature of 5GEx services, for instance the topology of value-added connectivity services impacts the number of intermediate providers concurring to deliver the service, and consequently its profitability margin for each of the involved providers.

An additional aspect analysed by 5GEx was the impact on prices of the model chosen to negotiate the service conditions. 5GEx analysed two possible options, auction model vs negotiation/e-negotiation model. Negotiation was deemed as better at the increase of the considered service's KPIs, and when the competition level is not high. E-negotiations are in general deemed better than direct negotiations, saving time and complexity, preventing verbal misunderstandings, and offering specialized tools to support the process.

4.1.5 Technical impacts

In general, the technical requirements are determined by and dependent upon the overarching collaboration tenets. The tenets bring to derive actual business procedures, demanding as minimum technical compliance level for federated members to support and implement a set of basic functions without whom it's not possible to carry on business along with federated peers. Certain technical components have to be implemented by each member, in compliance with the 5GEx technical design: e.g., the service repository, the SLA manager, and so forth. Besides such basic functions, a community (especially if closed) can decide to implement extra business logic, e.g. a penalty handling framework to manage SLA breaches, which in turn demand a deeper level of information sharing (e.g., deeper access to information generated by monitoring probes in the running providers' infrastructures).

In terms of impact on the technical design of the 5GEx architecture, the business model influenced some components of the architecture in the middle-lower layers. Nevertheless, the project chose to concentrate most of the features derived by the business analysis in a specific business layer named BSS/OSS, incorporating the dedicated functions and interfaces, with a clear-cut separation form the service layer embedded within the 5GEx orchestration components. The OSS/BSS layer was included in the architectural design, its implementation was instead deferred to follow-up heir projects.

4.2 Model 2 – Cloud federations

Research on Cloud Federations is at least ten years old and was driven by EU-funded research projects such as Reservoir [4] and Contrail [5]. At the same time some standardization work was carried on by IEEE and NIST working groups and resulted in the IEEE standard p2302, Standard for Intercloud Interoperability and Federation [6], and more recently in the NIST SP 500-332, Cloud Federation Reference Architecture [7], which defines general high level characteristics of cloud federations.

In this section, we summarize a classification of the federation approaches from several research projects and then highlight some pillars of the most recent NIST Cloud Federation standard.

4.2.6 EU-funded Research projects

Several EU-funded research projects studied cloud federation architectures. The major ones are OPTIMIS [8], Reservoir [4], Contrail [5], BonFIRE [9], mOSAIC [10], and EGI FedCloud [11]. Other researchers proposed further approaches, such as InterCloud [12], Cross-Cloud [13], Multi-Cloud [14], and Federated Cloud Management [15]. Kertesz, in [16], classified all the indicated works using four federation characteristics:

- *Hierarchical* federations, where a centralized management service coordinates the federation
- *Horizontal* federations, in which resources are exchanged among participating providers to optimize their utilization and reduce operating costs
- *Heterogeneity*, that indicates federations where participating providers can use different IaaS software stacks
- *Specialty*, a characteristic that indicated a unique capability of the proposed solution

Kertesz’s classification of the federation approaches proposed from the studied EU projects is summarized in the following Table 1.

Table 1 - Classification of federation approaches (from [16])

	Hierarchical	Horizontal	Heterogeneity	Specialty
OPTIMIS [8]	X	-	Yes	Legislation awareness
Reservoir [4]	-	X	No	Reservoir service stack
Contrail [5]	X	-	Yes	SLA contracts
BonFIRE [9]	-	X	Yes	Controlled networking
mOSAIC [10]	-	X	Yes	Cloud ontology, API
EGI FedCloud [11]	-	X	Yes	Virtualised EGI environments
InterCloud [12]	X	-	Yes	Market-oriented
Cross-Cloud [13]	-	X	Yes/No	Authentication
Multi-Cloud [14]	X	-	Yes	VM Mobility
FCM [15]	X	-	Yes	Meta-brokering

The ACCORDION federation, as it has been described in the Description of Action, can be indicated as Hierarchical, since it’s constituted by multiple Miniclouds coordinated by a central orchestration and optimization system, and is characterized by the Specialty of being oriented at federating Edge resources.

The EU Commission itself promoted a Cloud Computing Expert Working Group. In the last report about their findings [17], this Working Group identified several open research issues not fully covered by the previous research work. Among the main research issues, one is about Federation, Interoperability and Portability. The Working Group detailed several research tasks related to the Federation, such as application

management, scalable orchestration, networking and service composition. Another research issue identified is about Resource Management, with research tasks such as large-scale scheduling, heterogeneity, advanced reservation and energy efficiency.

4.2.7 NIST Cloud Federation Reference Architecture

According to the NIST Cloud Federation Reference Architecture [7] (CFRA), the essence of a federation is to "support the sharing of arbitrary resources, from arbitrary application domains with arbitrary consumer groups across multiple administrative domains".

The first and foremost requirement for sharing resources is to enable an authentication and authorization process across the organizations participating in the federation. This means that "Users in Organization A can discover and invoke services in Organization B, and Service Providers in Organization B can validate credentials from Organization A and make the proper access decisions", implying that trust relationships among all federation members must be established. The second NIST CFRA pillar is about sharing resources: federation members can share some of their resources and make them discoverable (based on a shared definition of their metadata) and accessible to other members. The third pillar is on the governance: federation members "agree upon the common goals and governance of their federation, based on well-known roles, attributes and policies".

Federation governance should include at least

- a process to grant or revoke membership
- a set of roles / attributes associated with the various actors
- a process to grant or revoke these roles or attributes

In more detail, the rules of a federation should define how the main federation lifecycle processes work. NIST CFRA identifies the main governance processes as follows: federation instantiation, membership, resource discovery, resource access, monitoring, reporting, accounting, auditing, incident response, termination.

5 ACCORDION federated business model

Starting from the baselines presented in the previous sections, the analysis performed by ACCORDION within its task 7.5 concluded that, although keeping into account other different inputs and references, there were quite strong similarities between the federation to draw up in ACCORDION and the one investigated within 5GEx. We identified multiple touchpoints between the two cases, coming from:

- the heterogeneity of treated resources (even if the resource/service hierarchy in 5GEx has a bit of extra complexity),
- the common nature of multi-provider/multi-administrative domain alliances,
- the presence of closely similar actors, like cloud providers, service providers and telecommunication operators, albeit in ACCORDION the scale of federated peers can generally be more oriented towards smaller providers, and the final product delivered by the federation is more inclined towards the application than the service scale

Therefore, the 5GEx work was taken as the main reference to define a set of initial business rationale to drive the ACCORDION federation. We kept off a number of items and complexities that are out of scope for ACCORDION, looking at capitalizing the most general aspects that could help us to sketch an initial business scenario, and offer some driving indications to the elaboration of ACCORDION architectural design. Particular attention was lent to the aspect of member on-boarding in the ACCORDION federation, an aspect potentially largely impacting the technical design and implementation.

The current deliverable is anyhow the first version of this techno-economic analysis of the federation. It will be updated and refined with its follow-up version, also taking advantage of the first step of ACCORDION validation. For instance, the federation on-boarding could be planned for the second implementation cycle, after an initial evaluation of the core functionalities in the ACCORDION system.

Federation openness

The 5GEx federation is substantially closed. The scale of involved actors and delivered services made it, in that scenario, too dangerous to open up admission for every potential participant, as well as defining an automated admission procedure, which is mandatory to enable fully open admission. On the other hand, the scale of participants made economically reasonable to accept the charge of costs requested to set up and enact a proper management of admission controls.

The ACCORDION case in this sense is a bit different. The scale of federation participants and the complexity of contractual agreements will very likely be reduced. This might hence reverse a bit the perspective, suggesting a more open approach to the federation model. Nevertheless, it is more careful to expect that some degree of pre-requisite offline negotiation might happen anyhow. A fully open admission would require an advertisement and entry point exposure of the federation, involving the implementation of automatic admittance and e-negotiation procedures that will not be dealt with in the first demonstration of the ACCORDION system. This point will be left open and eventually defined in the final version of the ACCORDION federation business model.

Federation entry point

Due to the predominant stance in the telecommunication operator realm, 5GEx chose to not implement a unique shared entry point, be it handled by one of the federation participants or by an external third-party brokering entity. In ACCORDION, this can be considered the initial choice as well. Hence, each of the federated entities will be a possible entry point to the ACCORDION system for new entrants. This sounds as a more reasonable first option, due to the expected highly distributed nature of the ACCORDION business federations, and to the possible granularity of participants. This choice can be re-evaluated for the final federation business model release, where the unique internal entry point option might be differently considered, especially for cases where one of the members (e.g., a telecommunication operator) have a significantly larger size and customer base compared to the others. The third-party broker is not expected to be a real option for the ACCORDION case, where it doesn't seem to be really justifiable.

Coordination model

The coordination model dictates how the federation governs the task of defining and provisioning the resource allocation for a given user request. In ACCORDION, it looks quite advisable to go towards a centralized model. The scale and typology of federation members suggests avoiding a distributed coordination, which would imply a significant amount of complexity for the providers, requiring them to install and run the whole ACCORDION core software at their premises. This does not look like a recommendable strategy in a federation where many peers are not so big, and have their resources mainly residing in the edge. It seems more efficient to centralize the ACCORDION core functionalities, like the aggregation and orchestration of resources in the continuum cloud-to-edge, and the service assurance capabilities, into a given central point. Thus, the candidate choice for ACCORDION is to have a fully centralized flavor of coordination model, where one of its members, potentially the one with a significantly different size intrinsically indicating it as an optimal access point, takes over the role of entry point for the external service requests. A telecommunication operator, also for its degree and ability of control on the network connectivity resources within the federation, might be a first option to take this specific role.

Governance model

In a federation like the ACCORDION one, a certain degree of openness looks almost inevitable, given the average size and fragmentation expected to characterize its members. A closed model, where a limited number of bigger and more “powerful” members implement a close control and a reduced trust kind of governance, does not appear to be ideal for enlarging the set of potential participants, and incentivizing the participation itself to this federation. The expected typical profile of ACCORDION federation members does not seem to pose a high risk of discriminatory practices or unfair pricing competition, such to call loud for tighter policies and a stronger level of control. Hence, in the first implementation we don't expect to implement special complex policies, nor sophisticated penalty charging mechanisms. These aspects will be more carefully analyzed in the final elaboration of the ACCORDION federation model.

Another aspect that will not be implemented in automated fashion during the first cycle is the mechanism of admittance and withdrawal. A more detailed analysis has been performed, nevertheless in the first implementation we will assume that the enforcement of federation joining agreements will occur offline, at

a business level rather than like an e-negotiation process or similar. In the initial implementation of the ACCORDION system, the only actual assumed constraint will be the need for the entrant member to deploy in its premises the piece of ACCORDION software required as a minimum to share resources in a cooperative way with other federated peers. The underlying intent is to promote the federation openness and foster its enlargement rather than making the access to federation functions more complex and cumbersome.

Provisioning model

In terms of provisioning model, ACCORDION is not making a defined choice sending prescriptive inputs to the technical tasks. In task 7.5, about this point it was decided to not take a decision at business model level, and adapt the final choice in the second round to the experience gained during the first implementation cycle. The ACCORDION federation might retain both the options (task forwarding vs capacity sharing), and the choice here can have a strong input from the observed technical complexity and implementation effectiveness.

Service assurance

The work of ACCORDION technology workpackages embeds at large extent capabilities to measure and store technical parameters related to the usage and performance of resources shared to deploy the client applications. ACCORDION design and implementation are well aware of the need, for the primary service point acting as one-stop-shop for the client application, to collect and evaluate the key technical metrics impacting the quality of experience for the client application, and to locate the critical points/resources compromising the overall QoE, or causing a breach of the committed service level agreements.

The final ACCORDION federation business model is expected to provide some indication on the shared service assurance mechanism. Based on past experiences, it was chosen to adapt the shared SLA handling to the actual mechanisms and functions technically implemented by ACCORDION, and not the other way around. An upfront business prescription here can easily turn out into relevant and not value adding technical complications, being the monitoring of resources not belonging to the same administrative domain an aspect which bears extra sensitivity issues. Ultimately, span of control in the ACCORDION federation will be refined based on the feedbacks from the first evaluation cycles, and those inputs will be highly helpful to nail down a realistic and sufficiently effective proposition of shared service level assurance.

Charging and pricing

Charging and pricing models are quite a complex part in the definition of a federated business model. Even a large project like 5GEx, after performing a very articulated and theoretically deepened analysis, did not eventually come up with a well-defined, unique and one-size-fits-all pricing scheme. The investigation of this aspect in ACCORDION did not select a final solution in its first round, preferring to have a feedback first from the technical implementation to better understand complexities and enacting tradeoffs. As already indicated, ACCORDION does not expect to have (at least in the first cycle) highly complex automation mechanisms for the contract agreement negotiation. At the same time, ACCORDION will implement reliable mechanisms to monitor the performance of resources employed by its client applications. Hence, as a minimal, an ex-post charging and pricing strategy is expected to be viable. In the first implementation, the complexity of

deploying and enacting automated mechanisms would not have been compliant with the need to have a first early checking of ACCORDION core functionalities.

Charging and pricing are also instrumental to address the achievement of KPI-O5-4, part of the ACCORDION objective 5 (Impact maximization) and touching upon the incentive for external providers/actors to become active members of an ACCORDION federation. This specific KPI aims at increasing the effectiveness resulting from the identified incentive measures for federation participation. A better description of how ACCORDION thinks to achieve this indicator will come in the final version of the current business model. This KPI does not exclusively look at economic incentives. The participation to an ACCORDION federation might have additional drivers, especially for mid-small size providers: e.g. increase of visibility on the market, opportunity to get in touch with other providers and establish additional business alliances/cooperations, optimization of own resource employment by associating and putting them in common with other providers, so reducing unproductive capital costs, and so forth. Nevertheless, an economic incentive is an ultimate driver to go, so it's important to show that this factor is present and has a tangible incentivizing effect.

To maximize as much as possible the economic incentive, you must act on two different sides: minimizing the costs of joining a federation, and maximizing the revenues stemming from such co-opetitive participation.

The latter point has a prominent business and marketing component: the more ACCORDION services are successfully delivered, the wider its customer base, the higher revenue level gets granted for all its participants. Further maximization of revenues can come from a possible diversification of business offerings. ACCORDION might look at differentiated service portfolios, encompassing higher priced services where for instance more valuable resources are provided, higher end quality of experience is granted, or more critical applications are served. On a technical view, this calls for the best possible capacity of the ACCORDION components, choosing and allocating resources to applications by making optimal choices, and handling the overall set of federated resources to serve as many requests as possible.

On the cost optimization side, there are both a general business effect and a general technical driver. The business effect is given by the opportunity to employ spare resources that the provider's own business has not been able to place on the market. At this purpose, ACCORDION is de facto extending each individual provider's addressable market, and anyhow the ACCORDION federation business model is not expected to be extremely compelling, rather leaving its members an important degree of control on the resources they are actually willing to share at any given moment. The reduction of unused capacity is a well-known, long-time standing problem for ICT service providers: a datacenter, and even a network, is typically oversized, being it dimensioned on the peak expected level of customer demand, which statistically only occurs for limited periods of time. A federation can help to at least partially fill up these resource employment voids, reducing unproductive capital costs and thus bringing a real benefit to profitability. Furthermore, in a particular market segment like the one of applications largely running in the edge, a federation can potentially allow to deliver services that a single provider could not offer, positively acting on the market dimension itself.

In terms of inputs for the technical design and realization, minimizing the cost of belonging to the federation means substantially making the technical provisioning of ACCORDION software into federation members as simple and cheap as possible. Demanding to have dedicated resources, technical and administrative, human

and technological, at an extent that can hinder or erode the gains achieved from being part of the federation business, should definitively be avoided. The ACCORDION design has been attentive to this aspect, looking, as much as possible in compliance with satisfying all its technical requirements, to not demand federated members the deployment of bulky software packages, complex to deploy and maintain, pursuing in first instance an architecture where the peer software may be a leaner local agent, easy enough to deploy and execute. The first round of implementation and evaluation will help to understand how much this wanted technical address proves to not be conflicting with the core technical capabilities that ACCORDION should provide.

Onboarding mechanisms

In the work performed by Task 7.5, a particular space and effort was devoted to analyzing the problem of members onboarding and offboarding, which might be differently dubbed as admittance and withdrawal of members to/from the federation. This aspect was considered particularly critical, and assigned a higher priority, since, from past experience, this was the one potentially more impacting technical design and implementation. Hence, an elaboration was carried on in tight cooperation with other tasks, first and foremost Task 3.5, to early detect possible critical issues, and anticipate functional aspects where a conflict between business rules and technical feasibility might have been potentially present.

To start this specific investigation, an initial set of naming assumptions was done to prevent possible ambiguities. In particular:

- We expressed the difference of the federation concept in business sense and in the technical sense. At business level, a federation is an association of resource providers willing to share their resources within an ACCORDION technical framework. In the technical sense, we intend a federation as a set of **Miniclouds** governed and controlled by a unique overarching ACCORDION orchestration entity.
- In turn, a Minicloud is intended as a pool of resources under control of a VIM (Virtual Infrastructure Manager), characterized by certain defining common parameters (e.g. capped latency)
- As **mechanism**, we mean an operational process implementing the functionalities and interactions among different components
- Finally, as **policies** we mean rules dictating the requirements to be met for enabling certain characteristics of the system, that can drive different aspects, e.g. security, eligibility, and so forth

Federating heterogeneous edge and cloud resources, which is the core concept to be realized by the ACCORDION technologies, implies the existence of an overarching association of infrastructure resource providers. Each of such providers owns some resources that can be made available to the federation, and participates to the process of adding/removing them to/from a given ACCORDION Minicloud. Defining a proper process for these operations of resource commissioning/decommissioning is crucial to enable a dynamic management of the ACCORDION pools, making the whole system far more innovative and powerful. Without an effective mechanism of resource onboarding, the system can only work in a static manner (i.e. with resource pools defined once and for all when a Minicloud is set up), losing key capabilities to adapt and guarantee a persistent service level anytime the context conditions change, or a resource providing node breaks down.

For the sake of simplicity, we can initially assume that the objects added and removed from the federation are identified with full physical nodes. Task 3.1 will define a resource data model, that will allow, in the second design cycle, to shift and lift the high-level analysis to a more granular resource level.

The key mechanisms to be defined, first in terms of business model, then in consequent technical implications, can be identified as:

- Joining a new node to the federation
- Decommissioning a node from the federation
- Verifying the actual presence and activity of a commissioned node

The diagram in Figure 6 below gives an overview of the main actions and flows that are part of this process:

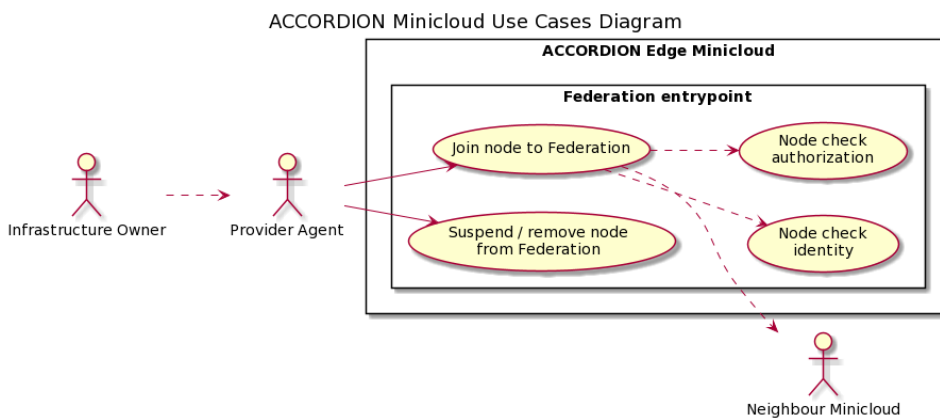


Figure 6 - Main functionalities for federation onboarding

In this analysis, we focused on the definition of onboarding mechanisms, not binding them to the type of federation admission policies. Hence, as anticipated few pages ago, we assume that all the nodes potentially participating to an ACCORDION Minicloud have a granted access to the federation, stemming from some pre-existing business agreement, in an ACCORDION federation expected to be substantially akin the open model. Furthermore, we assume that the Minicloud has already been established, and don't discuss here the Minicloud creation process, subject of a distinct investigation in Task 3.5 , as well as the criteria defining the definition of boundaries separating two different Miniclouds.

The main macro-distinction among the possible onboarding mechanisms is set by the definition of the actor in charge of initiating the process. We can have a **push-mode** onboarding (and off-boarding) mechanism, or a **pull-mode** one.

Push mode

In the push-mode mechanism, a node bearing resources is supposed to be aware of an existing ACCORDION Minicloud through which it could access the desired federation and initiate the process for joining it. This awareness can come from different ways, but it most likely follows up an offline business agreement taken between the infrastructure resource providers, opening up a dedicated endpoint to every node willing to join the federation. Following this agreement, the offering provider is supposed to install some software module

(e.g., an agent) provided by the ACCORDION federation, configured to enact a M2M joining procedure². The steps implementing this type of mechanism can be identified as:

- The joining node queries the Minicloud endpoint, that acts as an “access” point to the whole federation
- The Minicloud controller verifies the identity of the querying node, to authenticate it
- If the Minicloud security processes foresees it, there can be here an authorization step involved
- The access Minicloud, through verification steps of the proper parameters, identifies the best Minicloud where the node should join, and redirects to this latter the node’s admission request, to be completed there
- The final Minicloud controller returns an admission confirmation to the querying node, and asks for its resource descriptor
- The querying node provides its resource descriptor, and the endpoint(s) to access it over the control interface(s) foreseen by the ACCORDION architecture.

From this moment on, the node is part of the Minicloud resource pool, and can be requested to provision its available resources for any applications needing them. As well, it starts to communicate with the Minicloud monitoring component, according to the specifics defined by Task 3.1.

If the node provider decides to decommission its resources from the Minicloud, the decommissioning request can only be initiated at business level. The business agreement to be established should specify a “transition” interval, i.e. the width of a maximum time window after whose expiration the node is entitled to decommission its resources from the Minicloud. The key assumption is that, like in a “standard” public cloud, there is a capacity buffer available, ensuring that the resources of the quitting node can be migrated elsewhere without stopping the application, even if this latter is stateful.

Once the business agreement has been reached, an operational procedure starts with the following steps:

- The Minicloud controller is made aware of the need to remove the node from the federation (for good) or to suspend its subscription (temporary)
- The Minicloud controller tags the leaving node, to prevent any further assignment of the node resources to requesting applications
- The Minicloud controller checks if the node resources are currently assigned to and employed by one or more applications:
 - If the check is negative (resources unassigned), a decommissioning or suspending confirmation is returned to the node, and the resource descriptor is removed or flagged as inactive in the Minicloud resource repository
 - If the check is positive, the Minicloud controller spawns a hot-migration procedure, transferring the resources to a different available node in the Minicloud. Once the migration

² This process can also be replaced by a simple “manual” process, like the compilation of a web form spawning an initiation action from the Minicloud towards the joining node

procedure is completed, a decommissioning or suspending confirmation is returned to the node, as in the previous case.

- As further possibilities:
 - the node can be marked as "wants to leave" and not considered any more for scheduling new applications, when the current application exits the node is decommissioned;
 - the application components running on the node are migrated elsewhere and the node is decommissioned

During the working cycle of commissioned resources, the whole Minicloud system must enact a monitoring procedure, to ensure any anomaly is captured, first and foremost the occurrence of a node failure or disconnection from the Minicloud, and to ensure that resource usage is properly accounted for. Of course, also performance KPIs must be measured and any degradation detected, however these events don't ignite commissioning or decommissioning of nodes and are hence considered part of the general monitoring functionality.

In case of node failure undetected from the node itself (node breakdown) or node disconnection, the Minicloud controller must promptly react, putting the failing node in a "suspended" state and seeking other resources in the Minicloud to minimize the affected applications disruption. Once the applications have been recovered, the Minicloud controller keeps on polling the failing node status, ready to restore it as active node as soon as it gets up and running again.

If the node is able to self-detect a critical degradation of its resources, it proactively notifies the Minicloud controller, allowing it to enact a controlled resource migration in an as much as possible seamless way for the application users

Pull mode

In the pull-mode mechanism, the Minicloud controller is in charge of detecting candidate nodes to recruit and onboard into its managed federation. Even this model must necessarily subsume some previously offline agreement subscribed among the resource providers, as in the push mode case, since as a minimum the queried node must be able to understand and speak the ACCORDION control protocols. However, in this mode, the nodes don't proactively initiate their onboarding process, leaving it to be sparked only by the Minicloud controller. When the controller decides to post a request for a new node onboarding, it advertises its presence to the neighbor nodes on the network, asking receiving nodes to volunteer for making their resources available, optionally including a detail of the requested resource set. The controller should go through an advertising procedure, rather than directly inquiring a given node, since the availability status of node resources is not known to the controller in real time. Having the availability of such information would require an active cooperation of the nodes to keep it up to date, blurring the lines between the pull and push mode processes.

An alternative to the advertising process can be a cascading model. The Minicloud controller randomly picks a node from its existing list and queries it. If the query is successfully acknowledged, the node is onboarded. Otherwise, the node forwards the request to another one. This process is definitely more complex, implying

that each node in the range of the Minicloud is aware of the whole topology, or at least knows to whom neighbor node it can forward the request.

In case of pull mode, the decommissioning procedure is initiated by the Minicloud controller.

The monitoring procedure to handle node failures and disconnections is very much the same as in the push mode case.

Suggested choice

Ultimately, both push and pull mode require a previous offline agreement among the resource providers, otherwise the Minicloud controller and the federation nodes can simply not talk to each other with a common protocol and language. The push mode works well in a stand-by situation, where there is no immediate need of new resources, but a proactive onboarding increases the chances for the Minicloud controller to immediately find the needed resources when there is a new application to provision (or to upscale). On the other hand, the pull mode is more or less mandatory to manage the reactive cases, when the controller doesn't find the resources in its available resource registry, and must seek new nodes to add before having to reject the application's request. Bottom line, the onboarding protocol should include both capabilities: the pull mode handling is mandatory but adding the push mode can make the system more efficient and performant.

6 Inputs to technical design

Following the NIST model referenced in section 4.2.7, a Federation needs several governance processes to work well: federation instantiation, membership, resource discovery, resource access, monitoring, reporting, accounting, auditing, incident response, termination. Most of these processes should be supported by technical components and interfaces, both at the local level (i.e. resource provider, federation member, single site) and at the global federation level. This section tries to identify and propose such components and interfaces for each of the listed processes, following the NIST Reference Architecture.

6.1 Federation instantiation

Instantiation can be considered the birth of a federation. If it's started by only one entity, that entity can be considered the Federation Owner and Administrator; otherwise, if the federation starts through the signature of an agreement among some entities, the agreement itself must grant the role of Federation Administrator to one entity. In any case whole Governance processes should be well defined since day one. No specific technical components or interfaces are foreseen for federation instantiation.

6.2 Federation membership

The most important Governance process from a technical standpoint is about membership: it defines how resources can join or leave the federation pool. The assumption is that the entity owning those resources already signed a business agreement defining all the federation rules along with the rights and duties of each participant entity.

The joining resource should first contact a well-known federation endpoint, from which it receives a list of attributes to be valued (the equivalent of filling out a form) in a subsequent call. Such attributes include all the metadata that define that resource characteristics (to be defined by ACCORDION Task 3.1), the admin credentials to access and operate on the resource and some data identifying the resource itself (at least its IP address). The joining resource may also receive a script to be executed locally to check some prerequisites and determine the closest (in terms of latency) federation entry point to be contacted for the actual join operation. The provided script can even help in determining the needed values for the required metadata.

The actual "join federation" call will pass as parameters all the values for the indicated attributes and should be signed with the private key of the joining entity to both ensure about the source and the integrity of the provided information. In the first release of the federation services it will be acceptable an implementation that takes full control of the resource; whereas further releases should implement the possibility of partial commitment, to allow the provider entity to retain some control on its resources: this should be better analysed in the next year. The ACCORDION component implementing the two methods indicated above can be named Federation Endpoint and be hosted in each Minicloud that is part of the federation.

The "join federation" method, after verifying the signature, will check the provided information and in case it is valid will first answer success to the caller and then will use the provided credentials to install all the

required ACCORDION software on the joining node and will publish the resource descriptor in the Minicloud Resource Repository, marking it as “available”.

The same Federation Entrypoint component should also implement a “leave federation” method, which first marks the resource as “leaving” in the Minicloud Resource Repository and then checks if the resource is currently allocated to run some application. If the check is negative, a decommissioning or suspending confirmation is returned to the caller, and the resource descriptor is removed or flagged as inactive in the Resource Repository. If the check is positive, the request is rejected and could be repeated later; in the meantime the resource will continue executing the running application, but no further jobs can be scheduled on it because the resource has been marked as “leaving”. In case of a leaving resource, the scheduler can also optionally prepare to migrate the application to a different node or even to a different Minicloud.

6.3 Resource discovery

Resources that are part of the federation are first registered in the local Resource Repository of the Minicloud through which they joined the federation. But the resources should be available to the whole federation, therefore the Minicloud Resource Repository is expected to either publish the local resources to a global federation-wide repository, or at least index the local resources in a way that could be discovered globally. Resource Indexing and Discovery is a responsibility of ACCORDION Task 3.2. Here it’s just worth pointing out that some local Resource Repository should be available in each Minicloud and it should be possible to mark each resource with a status which can be at least “available”, “leaving” and “unavailable”. Another useful state could be “loaded”, possibly paired with a dynamically updated value representing the resource load. It should be possible for the federation-level ACCORDION orchestrator to query the Resource Repository of each Minicloud, in order to understand which resources can be used to run a given application. Resources with status “leaving” or “unavailable” must not be considered for scheduling new applications. The possibility for the providers to define specific discovery policies for their resources, for example to avoid hosting certain types of applications or users from certain domains, can be researched in the next year. The resource states proposed in this subsection are summarized in the following State Diagram.




Figure 7 - ACCORDION Resource State Diagram

6.4 Resource access

When the resources are available in the ACCORDION federation they can be allocated for running user applications or to be directly accessed by users. In ACCORDION, only application owners or developers are

considered as possible resource users, whereas the end users of those applications are not considered and their identification and management are a responsibility of the applications themselves.

Some Federated Identity system must be available to authenticate both users that need to access federation resources and providers that offer them. Moreover, there should be some way (e.g. a Policy Enforcement Point for each Minicloud) to check the authorization of users to access the resources, either based on centrally defined roles (RBAC) or based on access policies defined by the single providers / Minicloud administrators, or a combination of both.

6.5 Monitoring, Reporting, Accounting, Auditing

In a complex and potentially large system such as a federation, it's vital to collect all the information needed to keep the system running, to keep it secure, to optimize it and to satisfy all the expectations of both users (e.g. QoS, QoE) and providers (e.g. profit).

Monitoring of both resources and applications is a responsibility of ACCORDION Task 3.1. Here it's anyway needed to point out that the type of information collected may be different for each of the goals indicated above, and both the collection goal and the type of information can influence the collection frequency, the recording and distribution medium, the aggregation need, archiving and aging policies, and even access policies. For example, availability and performance information should be collected with a relatively high frequency, but may not need to be kept for a long time, whereas usage and access information may be collected in the form of events, but should be kept for longer to satisfy both accounting and security (even possibly forensics) requirements. It's possible that such different requirements may be satisfied by different ACCORDION components. For example, Monitoring components can collect high frequency and volatile information; dedicated Filtering or Aggregator components may be designed to analyse, compose or correlate different information or information from different sources; Event Recording and Logging components may be dedicated to handle information that satisfies accounting and auditing needs. In the end, several components can be added to the architecture depending on the goals to be satisfied.

One further point to be considered is that Monitoring, Aggregation and Logging components generate new data that must be protected: access to monitoring information should be regulated by specific access policies to guarantee both confidentiality and fair competition among providers.

6.6 Incident response

The identification of incidents will be done by dedicated Filtering components, processing data produced either at Minicloud or at federation level by Monitoring, Event Recording and Logging components. Different types of incidents can be identified and, for each type, a different Filtering or Aggregator component is needed. To identify QoS violations (e.g. if the availability of a node falls below the promised threshold) a filter will be needed on Minicloud local monitoring data. To identify QoE violations (e.g. the application end user is distracted by a lagging UI) some aggregation and processing on monitoring data at both Minicloud and federation level will be needed to evaluate a QoE model (such as that produced by ACCORDION Task 6.3).

The identification of security violations will be possible by using filtering and correlation components applied to event logs, both locally and globally. Specialized analyser components applied to the same event logs can even identify frauds.

The response to identified incidents will depend on the incident type. Replication and migration techniques can be used (it will be decided by ACCORDION Task 4.3) to mitigate QoS and QoE violations. Stricter policies may be added to cope with security or fraud incidents.

6.7 Termination

The NIST standard also considers the possibility that a federation terminates its existence, for example when it's not being used or when it's no more economically viable. This case will not be analysed in ACCORDION, but a procedure for orderly termination should be agreed anyway as part of the governance processes for any federation.

6.8 Summary

The components identified in the previous subsections and their relationships are shown in the following UML diagram. Please note that the identified set of components is not an indication of what will be implemented, but only a suggestion for the architecture definition of the next project iteration.

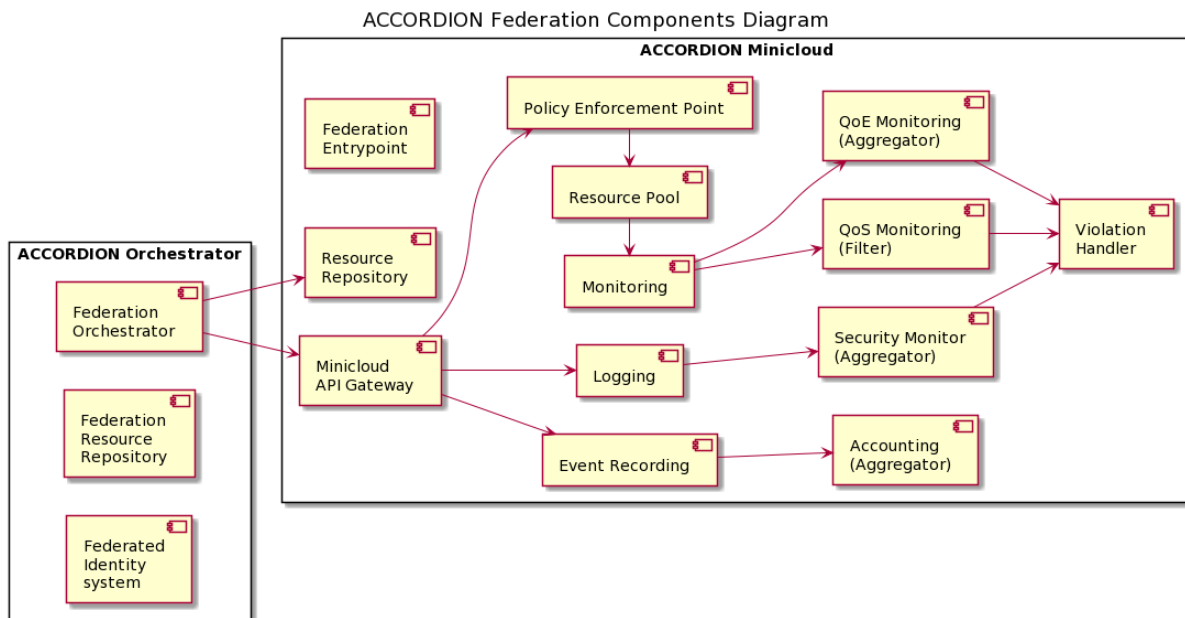


Figure 8 - ACCORDION Federation Components Diagram

7 Conclusions

This document is the result of the analysis work done in ACCORDION Task 7.5 on federation models and federation business logic. The document analyses different federation models in terms of their structure and governance, admittance and coordination policies. The main models analysed are the model proposed by the FP7 5GEx project, a federation of 5G telecom providers, the cloud federation models proposed by several EU research projects, and the NIST Cloud Federation Reference Architecture.

Based on the previous analysis, a model suitable for the ACCORDION Edge federation is proposed and analysed. From the proposed model and from the NIST Architecture the analysis derives some technical suggestions for the identification of architectural components for ACCORDION.

The current deliverable is the first version of this techno-economic analysis of the federation. It will be updated and refined with a follow-up version, adding further analysis and taking advantage of the first step of ACCORDION validation. A couple of important federation aspects that will be analysed in the next deliverable are the pricing model for the ACCORDION federation and the incentives to obtain that more providers join the federation. A first way of supporting the enlargement of the edge infrastructure pool has been to select an open and lightweight federation model, i.e. one that doesn't require complex procedures and many resources to be managed. Another aspect that will be more carefully analyzed in the final elaboration of the ACCORDION federation model is the need to implement additional controls and special policies to prevent discriminatory practices or unfair pricing competition.

8 References

- [1] *Encyclopedia of Cloud Computing*, S. Murugesan and I. Bojanova, eds., Wiley-IEEE Press, 2016.
- [2] A.N. Toosi, R.N. Calheiros, and R. Buyya, "Interconnected Cloud Computing Environments: Challenges, Taxonomy, and Survey," *ACM Computing Surveys*, vol. 47, no. 1, 2014, pp. 7:1–7:47.
- [3] D. Kogias, M. Xevgenis and C. Patrikakis, "Cloud Federation and the Evolution of Cloud Computing" in *Computer*, vol. 49, no. 11, pp. 96-99, 2016.
- [4] Wolfsthal, Y., Elmroth, E., Ben-Yehuda, M., & Emmerich, W. (2009). The reservoir model and architecture for open federated cloud computing. *Journal of Research and Development*, 53(4), 1-17.
- [5] Cascella, R. G., Blasi, L., Jegou, Y., Coppola, M., & Morin, C. (2013, May). Contrail: Distributed application deployment under SLA in federated heterogeneous clouds. In *The Future Internet Assembly* (pp. 91-103). Springer, Berlin, Heidelberg.
- [6] IEEE standard p2302, Standard for Intercloud Interoperability and Federation (SIIF), see <https://standards.ieee.org/project/2302.html>
- [7] NIST SP 500-332, Cloud Federation Reference Architecture, online at <https://www.nist.gov/publications/nist-cloud-federation-reference-architecture>
- [8] Ferrer, A. J., Hernández, F., Tordsson, J., Elmroth, E., Ali-Eldin, A., Zsigri, C., ... & Ziegler, W. (2012). OPTIMIS: A holistic approach to cloud service provisioning. *Future Generation Computer Systems*, 28(1), 66-77.
- [9] Jofre, J., Velayos, C., Landi, G., Giertych, M., Hume, A. C., Francis, G., & Oton, A. V. (2014). Federation of the BonFIRE multi-cloud infrastructure with networking facilities. *Computer Networks*, 61, 184-196.
- [10] Petcu, D., Di Martino, B., Venticinque, S., Rak, M., Máhr, T., Lopez, G. E., ... & Stankovski, V. (2013). Experiences in building a mOSAIC of clouds. *Journal of Cloud Computing: Advances, Systems and Applications*, 2(1), 12.
- [11] EGI Federated Cloud, see <https://egi-federated-cloud.readthedocs.io/en/latest/index.html>
- [12] Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., & Brandic, I. (2009). Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation computer systems*, 25(6), 599-616.
- [13] Celesti, A., Tusa, F., Villari, M., & Puliafito, A. (2010, July). How to enhance cloud architectures to enable cross-federation. In *2010 IEEE 3rd international conference on cloud computing* (pp. 337-345). IEEE.
- [14] Bernstein, D., Ludvigson, E., Sankar, K., Diamond, S., & Morrow, M. (2009, May). Blueprint for the intercloud-protocols and formats for cloud computing interoperability. In *2009 fourth international conference on Internet and web applications and services* (pp. 328-336). IEEE.
- [15] Marosi, A. C., Kecskemeti, G., Kertesz, A., & Kacsuk, P. (2011). FCM: an architecture for integrating IaaS cloud systems. IARIA.
- [16] Kertesz, A. (2014). Characterizing cloud federation approaches. In *Cloud Computing* (pp. 277-296). Springer, Cham.
- [17] Schubert, L., & Jeffery, K. (2012). Advances in clouds. *Report of the Cloud Computing Expert Working Group. European Commission*. Online at <http://www8.cs.umu.se/kurser/5DV095/VT13/literature/schubert2012advances.pdf>
- [18] 5GEX project, Deliverable 2.3: 5GEx Business and Economic Layer