

**ECMWF/ItpQ/2020/307**

**Invitation to pre-qualify for "European Weather Cloud – Delivery and Operation of a Production On-Prem OpenStack Cloud"**

**Clarifications V3**

**Issued: 11 February 2021**

ECMWF/ItPQ/2020/307 Questions and Answers

Number	Question	Answer	Published
1	<p>Could you provide more detail on what content is expected on the Standards&amp;Procedures section? We read in the ItPQ "Responders are encouraged to give any supporting Standards &amp; Procedures information that they feel lends weight to the views expressed in their submissions." however it is not fully clear to me if the ECWMF expects international standards or best practices, GMV methodologies, standards and/or procedures in this kind of services, specific technical procedures for the solution or something else.</p>	<p>Responders are, of course, invited to offer any and all Standards &amp; Procedures that they feel lend weight to the views that they express in their submissions. The S&amp;Ps that are most likely to attract positive attention from ECMWF are those that are directed at the technology and operational environment that they propose in their submission. Perhaps, to this end, for instance, responders may like to describe in brief the relevance of those S&amp;Ps that they offer.</p>	13-Jan-21
2	<p>Is there any possibility to participate in the ItPQ as a consortium of two or more companies, that would jointly prove their competence? If not is there a possibility to use subcontractors and prove the competence using such subcontractors.</p>	<p>ECMWF will only enter into a contract with a single legal entity. We recognise that some responders may wish to involve and work together with a number of organisations to deliver the requirements. In such cases, these organisations must identify a lead contractor who will sign the contract with ECMWF and who will be responsible for putting in place legal arrangements to ensure that it can guarantee that all other organisations will also meet the contract obligations. This can be in the form of a subcontracting. ECMWF will not enter into multiple contracts with individual members of consortia or groups of service providers.</p> <p>In this regard, suppliers should note that partial responses, by suppliers who are interested in or only capable of offering complementary services or solutions, will not be considered.</p>	13-Jan-21
3	<p>In order to have a better understanding of the context, please indicate where ECMWF intends to deploy the new OpenStack solution: if in Bologna at the new Data Center as seems to be evident from the documentation, or in Reading at the current Data Center or in distributed mode among the two Data Centers.</p>	<p>ECMWF plans to make all deployments under the eventual (prospective) contract in its new data centre in Bologna. There are no plans to make any such deployments in the existing data centre in Reading</p>	13-Jan-21
4	<p>Section 4.3.3 of the document is titled "Three-Nines Planned Downtime (MHLR#3).", however the accompanying text describes two-nines availability (and this is referenced elsewhere in the document). Please can you confirm how many nines availability?</p>	<p>Two 9s (heading is wrong, text is correct).</p>	28-Jan-21
5	<p>4.3.1 states that Ussuri is the required API version, is this due to feature availability on the API or down to length of support?</p>	<p>Both feature baseline and length of support. Note that you may argue the case for later than Ussuri, but not earlier than Ussuri.</p>	28-Jan-21
6	<p>4.3.9 states that some users may require access to bare metal instances, but no mention is made regarding Ironic being a required component. Are the bare metal instances expected to be deployed by OpenStack, or is there just a need for connectivity between OpenStack and these bare metal servers?</p>	<p>In respect of section 4.3.9 this implies the use of Ironic.</p>	28-Jan-21
7	<p>4.3.1 states that The base OS must be either CentOS 8 (CentOS 7 or earlier is not acceptable) or Ubuntu 20.04 LTS (18.04 LTS or earlier or any non-LTS release is not acceptable). Assuming that a proposal based upon RHOSP is acceptable, would it also be acceptable to propose RHEL 8 as the O/S?</p>	<p>You may propose RHEL if you so wish and provided that you state your reasons for doing so in your response. Please note that as per Section 6.4.3. "In such cases, license fees shall be made clear in terms of licensing term, licensing volume and commercial support.</p>	28-Jan-21

ECMWF/ItPQ/2020/307 Questions and Answers

8	1.7.3 - Please note that ECMWF has a strong preference for not using proprietary licensed software or software that otherwise carries a fee for use, whether one time or subscription based. Whilst we understand the rationale behind this statement in terms of cost reduction, as a commercial provider of public and private clouds our standard approach is to deploy software that benefits from a defined support framework (i.e. commercially provided software). For this reason would an Open Source solution that is backed by a commercial support agreement be an acceptable alternative to Open Source community based support?	ECMWF is prepared to consider such a support arrangement. Please be reminded though that in Section 1.1 we state that "price will carry a significant weight during any subsequent ItT evaluation."	28-Jan-21
9	3.1 states (figure 1) that Magnum, Zun and Quinling are mandatory requirements. Would the customer consider alternatives to these components that deliver comparable functionality?	ECMWF is prepared to consider alternatives to Magnum, Zun and Quinling; however, responders should be aware that we reserve the right to find the offered alternative non-compliant with our requirements.	28-Jan-21
10	4.3.12 - With regards to CMDB services provided by ECMWF, what access if any would a provider have to update/modify and use data contained within the CMDB. For example automation tooling may require information stored in the ECMWF CMDB, would a provider have access to this, or would the provider be expected to maintain their own CMDB for that purpose.	ECMWF envisages this being a joint development effort so as to provide tight integration and to avoid replicated / disconnected systems. Where possible the vendor should interface via APIs with the existing instances of OpsView and NetZoom deployed at ECMWF.	28-Jan-21
11	6.3 - The bid asks for costs of network switches (data and out of band), racks, switch gear and power distribution - will these not be provisioned as standard in the new Bologna data halls, or do they need to be provided specifically for this project?	The new Data Centre will deliver leaf-switch ports as required for interconnection with the OpenStack environment. Switches internal to OpenStack or at the edge of OpenStack will need to be provided as a part of the project. Note that the OpenStack environment will never have SDN control of any of the enterprise DC network switches. Racks and power distribution units internal to the racks also need to be provided by the vendor.	28-Jan-21
12	In order to clarify the availability required for planned downtime, because in the section header an availability of three-nines is specified while in the section description an availability of two-nines is established, what would be the availability required on the platform when performing interventions that are within this scenario?	Two 9s (heading is wrong, text is correct).	28-Jan-21
13	In order to better understanding of the environment, in the section for the selection of the Kubernetes distribution and its deployment method, reference is made to Kubespray and Magnum which can be considered as deployment methods for Kubernetes clusters. In the case of Rancher, do you want to compare only the use of RKE as a deployment tool? Or do you want to have a global vision of the use of Rancher as a management tool for various Kubernetes clusters?	ECMWF is keen to acquire the broadest possible understanding of the market; to this end, we advise you to describe the use of RKE for deployment and to describe the use of Rancher as a management tool for various k8s clusters.	28-Jan-21
14	In order to better understanding of the environment, there is a reference to a multi-tenant Kubernetes cluster hosted on bare-metal servers. Is there an estimate of the number of nodes that will form this cluster? Are these servers additional to those established for the deployment of OpenStack? Does the deployment and administration of this cluster have to be done with the same tool as the single-tenant clusters deployed in OpenStack?	ECMWF requires to be able to scale its server farm to some 500 servers. This farm will be used to provision OpenStack, K8S (on bare metal) and possibly other requirements. Any limitations that we should be aware of in respect of clusters sizes should be stated in responses. It is acceptable to propose one or more tools as responders see fit to meet the multi-tenant and single-tenant cases.	28-Jan-21

ECMWF/ItpQ/2020/307 Questions and Answers

15	In order to better understanding of the environment, in relation to the handling of identity access, will this access to the different Kubernetes clusters be integrated with the global identity management service? Or will access to single tenant clusters be done on a standalone basis?	Single tenant clusters will be standalone and/or integrated with and managed by Member States / Partner security services.	28-Jan-21
16	It is understood that ECMWF values stability, security & roadmap reliability/sustainability as crucial assets of forthcoming EWC ecosystem. As per MHLR#1, bleeding-edge solutions shall not be considered. As a supplier and a bidder we are in the position to inform that Qinling is largely an abandoned project with no valuable commits for the last 2 years, planned to be EOLd in the OpenStack Wallaby. Thus in our recognition it meets the criteria to fall into the 'bleeding-edge' category. To better understand & accomodate the requirement, we would like to understand what is the exact use case behind Qinling that perhaps can be met by other solutions?	The requirement is to be able to support 'function as a service' on the OpenStack infrastructure. ECMWF thanks you for following the advice in Section 3.6 - namely "If a requirement is thought to be unsound or otherwise misguided, responders are advised to bring this to the attention of ECMWF during the response preparation phase". ECMWF will not use QINLING as compliance criteria; however, all suppliers in a position to do so are encouraged to provide advice on its use.	28-Jan-21
17	<p>It is understood that ECMWF values stability, security &amp; roadmap reliability/sustainability as crucial assets of forthcoming EWC ecosystem. As per MHLR#1, bleeding-edge solutions shall not be considered. As a supplier and a bidder we are in the position to inform that Zun development has been practically put on hold recently and it comes with a number of serious bugs and shortcomings. To name a few:</p> <ul style="list-style-type: none"> <li>• Instability of the service itself</li> <li>• Poor support for Python 3.x</li> <li>• 'Disappearing' networks</li> <li>• Placement service issues</li> <li>• Poses security and stability risk to the undercloud as both Zun and OpenStack control plane run on the same Docker instance (in case of a bug, it can freeze or compromise control plane gaining root access)</li> <li>• Lack of container orchestration means also puts to question its use in the favor of K8s</li> </ul> <p>All of the above disqualifies Zun (and associated project Kuryr - that shares Zun's downsides) from production use. To better understand &amp; accomodate the requirement, we would like to understand what is the exact use case here, so that we can suggest an alternative approach.</p>	The use case is "quick and simple testing of small numbers of containers using the Cloud infrastructure rather than a local workstation". ECMWF thanks you for following the advice in Section 3.6 - namely "If a requirement is thought to be unsound or otherwise misguided, responders are advised to bring this to the attention of ECMWF during the response preparation phase". ECMWF will not use ZUN/KURYR as compliance criteria; however, all suppliers in a position to do so are encouraged to provide advice on their use.	28-Jan-21
18	<p>MHLR #8</p> <p>Can you please describe the use cases for IPv6 in EWC? (e.g. is it mandatory in backend infrastructure, control plane of OpenStack and Ceph, cloud virtual infrastructure (VMs and networks), cloud Internet access via IPv6)?</p>	ECMWF would expect to be able, in due course, to exclusively use IPv6 in all aspects of its Data Centre and external connectivity settings.	28-Jan-21

ECMWF/ItPQ/2020/307 Questions and Answers

19	<p>MHLR #16 Will EWC be located at the edge of ECMWF network (and so it is assumed security side of the project shall be furnished by the supplier) or it will pass through standard network security measures deployed by ECMWF? If the latter, which ones?</p>	<p>In accordance with the Bologna network high-level design, infrastructures that are partially or totally managed by external entities will be deployed in a network that is segregated from the rest of the LAN within a specific security zone and behind the data centre firewall, which means that it will benefit from the network security controls provided by ECMWF (IPS, SIEM...). In terms of operations security, we are applying the defence in depth design principal. The relevant stakeholders will need to implement the required security controls and comply with ECMWF's security policies (e.g.: password management and technical vulnerability management). The exact HLD and LLD designs and the relevant network security controls will need to be discussed and agreed with all relevant stakeholders; however, this is a matter to be elaborated on at the ItT stage and beyond.</p>	28-Jan-21
20	<p>MHLR #16 Will the tenant (Member States, Partners) access to EWC be provided through the plain Internet or any methods of narrowing this access at the network level is planned? (VPN, tunnelling, firewall, ACLs, others?)</p>	<p>ECMWF does not have current plans for Member State / Partner access above and beyond SSH and/or HTTPS and/or other secure bespoke protocols. However, responders in a position to provide Virtual Private Cloud functionality should consider offering it.</p>	28-Jan-21
21	<p>MHLR #17 Can you please elaborate more on envisioned guaranteed levels and ring-fencing scheme on compute, storage and network levels? Is it required to guarantee exclusivity of usage of compute and storage space or also reservation is required? If the latter, does it need to happen specifically on a domain or project level? Does it need to be timed (expiring leases) or permanent (reserved until feel differently) What is the exact business model and use case?</p>	<p>ECMWF Member States and Partners may wish to place hardware that they have paid for in the ECMWF data centre. They would then require full and exclusive use of that hardware. Such hardware would be used in some mix of the following roles (not necessarily exclusive): compute servers, network servers, ceph OSDs, complete ceph clusters. Responders are encouraged to describe the possibilities and limitations of such ring-fencing.</p>	28-Jan-21
22	<p>MHLR #17 What kind of integration of EWC resources with current Member States/Partners environments is needed?(VPN site-2-site setup in a (semi)automated way, a full VPNaaS or others?)</p>	<p>ECMWF does not have current plans for Member State / Partner access above and beyond SSH and/or HTTPS and/or other secure bespoke protocols. However, responders in a position to provide Virtual Private Cloud functionality should consider offering it.</p>	28-Jan-21
23	<p>ItPQ says that ECMWF expects to provide at minimum, compute server local ephemeral storage, Ceph-based HDD storage and Ceph-based SSD storage. Table 3 requires 40 PB for Hard Disk. Could you please specify how this 40 PB is distributed? This is, the quantity for compute server local ephemeral storage and the quantity for Ceph-based HDD</p>	<p>As a guide, ECMWF would suggest that the ephemeral storage be provisioned in reasonable proportion to the amount of memory distributed across all the compute servers.</p> <p>HDD and SSD for all other storage needs should be provisioned in inverse relationship to the costs per unit-of-storage of such devices. This guidance is necessarily less stringent at this stage; however, it may become more so as hard use case data becomes available.</p> <p>Responders are advised to price all storage types separately for maximum flexibility during evaluation of responses.</p>	28-Jan-21

ECMWF/ItPQ/2020/307 Questions and Answers

24	In case of the base OS of the platform, CentOS, will no longer be updated (like in December of this year) and will become an upstream branch called: CentOS Stream. Are you consider using CentOS Stream as a base OS? (Page 22 - Section 4.3.1)	You are advised to state your reasons for your choice of underlying operating system. Given the change in status of CentOS announced by RedHat since the issue of this ItPQ, ECMWF will not disqualify a responder on the basis of a non-compliant OS choice. However, you must provide sound reasoning for your choice. And, we cannot guarantee that we will allow such a choice at the ItT stage, by which time we expect the licensing situation concerning RHEL/CentOS will have become clearer. See also Q7 above.	11-Feb-21
25	Regarding the required storage capacity described in table 3 (page 21 - section 4.2), it refers to the available storage or the total that will be distributed under a certain type of protection (i.e. RAID 1, RAID 5, etc).	All sizes refer to USABLE storage, NOT aggregate raw storage.	11-Feb-21
26	Regarding the number of servers described in table 3 (page 21 - section 4.2), do these 500 servers include management, control, computing, etc. or only computing?	Servers of all types including bare metal. Not just OpenStack compute servers.	11-Feb-21
27	For the Disaster Recovery (page 23 - section 4.3.4) we would like to consider the use of a European cloud provider as a shelter when applying DR, is this approach acceptable to you?	This section is intended primarily to address the policies, procedures and mechanisms required to achieve 2*9s DR ability. Bologna DC has two segregated data halls and this, for the time being, is considered sufficient physical capacity to achieve our aims.	11-Feb-21
28	Does the spine and leaf infrastructure in Bologna have any limitations? Does the DC in Bologna include Racks and the TOR network, or do we have to set it up?	The OpenStack environment will be self-contained. Racks and TOR switches must be supplied and configured according to Bologna DC interfacing standards that will be provided during the ItT phase of the project. Please also see Q19 above.	11-Feb-21
29	In the case of using the GEANT network, is there the possibility of receiving users from the internet to the platform? If positive, they can load data and how much do you estimate this load to be?	We expect a number of users to access via the Internet without distinction as to whether these users are routed in via GEANT, national NRENs, or any other ISP. At present there is no estimate on user traffic patterns however please note that the Bologna Data Centre will be connected to the Internet using 100Gbps links and users may need to upload large meteorological datasets , in the order of Terabytes, in a timely manner.	11-Feb-21
30	The IOPS and durability requirements have not been posted for the HDDs/SSDs and therefore it's difficult to balance cost vs performance. For example the largest mixed-use SSD is limited to 1.92TB while Read-Intensive drives exceed 3.5TB among other interface considerations like SATA vs NVME. Can further detail be provided to enable cost-effective scoping please?	Please quote the range of cost/performance data for all readily available devices. As more use case data becomes available we expect to be able to refine this requirement; but this will not be possible during the ItPQ phase of the project.	11-Feb-21
31	ECMWF has requested 'an indication of an implementation budget with breakdown for each of the features presented and an annual budget for all commercial licences.' Although it is clear ECMWF is open to considering other options, is the preference to avoid a commercial construct where hardware, software and services are delivered as a monthly cost for a defined duration?	ECMWF is willing to consider such an arrangements though it has a requirement to build up its in-house capabilities. Please ensure that full lifecycle costs are covered and that full technology disclosure is possible.	11-Feb-21

ECMWF/ItPQ/2020/307 Questions and Answers

32	Following on from the above, would a proposed support mechanism akin to a service outsource of some and/or all of the in-scope requirements be non-compliant?	ECMWF is willing to consider such a model. It is important that the service is delivered on a collaborative basis - full access to all the technology by ECMWF is required. Please note the capability goal stated in section 3.5 of the ItPQ.	11-Feb-21
33	Is it possible to share a high-level support RACI to better understand the areas of the collaborative delivery approach ECMWF is looking for suppliers to own? For example, monitoring at each layer (network vs. Infrastructure vs. software), break-fix, spare and capacity management, maintenance, upgrades etc.	Devising such a responsibility assignment is one of the objectives of the ItPQ exercise.	11-Feb-21
34	Please can ECMWF elaborate on 'whether on-payroll or on a direct contract or on a support contract' in respect of support staff? Is this referring to ECMWF's payroll (as opposed to a supplier's) and if so is it envisaged that suppliers will interface into ECMWF's existing SIAM process as a hand-off partner for second and third-line support?	This section of the document refers to the ECMWF payroll. The wording should be taken to mean ECMWF permanent staff as opposed to other ECMWF resources.	11-Feb-21
35	ECMWF have clarified availability requirements but it is not clear whether this is a KPI or an SLA with an envisaged service credit mechanism. Please can this be confirmed?	ECMWF sees it as likely that initially we will monitor KPIs. As experience builds, we may require to move to some level of SLA.	11-Feb-21
36	Currently OpenStack loadbalancer does not support load balancer members affinity across Data Halls (computes only), which in turn implies usage of other load balancer solution external to both Data Halls. The load balancer in question will not support OpenStack API though. Is it acceptable? If not, can you please describe the exact use case, expected failover time and other SLOs? If loadbalancer solution is required, please advise who will be responsible for its delivery (ECMWF/vendor).	At this stage, we require Octavia within each OpenStack cluster. You are encouraged to describe the options for achieving cross-data-hall load-balancing.	11-Feb-21
37	Are there any DCI bandwidth limitations for (DH1-DH2) traffic or (preferably) vendor can specify own bandwidth requirements to be delivered by ECMWF as a prerequisite CFI.	Cross data-hall spine capacity is 400Gbps shared among all ECMWF workloads including the HPCs.	11-Feb-21
38	Does ECMWF foresee a need for private loadbalancers in OpenStack tenant (Octavia) with members in separate DHs (be it AZ, Regions..)? Using vanilla OpenStack Octavia load-balancer can distribute traffic within single Availability Zone only. Does ECMWF find this solution acceptable?	Please see answer to 36	11-Feb-21
39	Is L3 connectivity in tenant's virtual private network between clouds in DH1/DH2 required? Currently OpenStack imposes certain limitations on virtual networks – they cannot span over multiple Data Halls, therefore comms is available via external networks with all their implications stemming from particular use case scenarios that need to be discussed in detail.	We expect to work within the limitations of the current versions of OpenStack. In due course, new networking functionality may be adopted if sufficiently attractive for future use cases.	11-Feb-21
40	Would ECMWF accept a solution where the communication between servers in the following relations would only be done over "External" networks? 1. Bare Metal in DHx – Bare Metal in DHy 2. Bare Metal – VM 3. Bare Metal in DHx – VM in DHy Can ECMWF describe exact use cases?	External network connectivity of these three cases is acceptable at this time.	11-Feb-21

ECMWF/ItpQ/2020/307 Questions and Answers

41	In our understanding, the initial approach to the EWC would be 500 computes divided across two Data Halls. Shall these be deployed in this quantity as two clouds, (one per Data Hall) or a collection of smaller clouds with a common identity management?	There are likely to be a minimum of two clouds active in production all of the time - one in each data hall. There would then be several smaller clouds instanced as required by various development, test and pre-production activities.	11-Feb-21
42	4.3.1 – In response to a previous CQ you have stated that “you may argue the case for later than Ussuri but not earlier than Ussuri”. For operational and support reasons any bid we submit will be based upon Red Hat OpenStack. The current release (RHOSP16) is based upon the Train OpenStack code base. Red Hat release a new major upgrade every third community release, so the next major upgrade (RHOSP17) will be to Wallaby, however, in H2 2021 key features from Ussuri and Victoria are planned to be released as intermediate updates to OSP16, and will therefore be available prior to the commencement of contract. Will a submission utilising RHOSP16 be deemed non-compliant given that it does not have the required Ussuri?	We will not deem a response offering RHOSP16 non-compliant provided a clear roadmap to Wallaby and beyond is included in the response.	11-Feb-21
43	You mentioned that GPUs wil form part of the requirement. Do you have a feel for what percentage of the servers GPUs will be required. Also are they used to accelerate workloads or more for ML type activities?	We do not yet have a good feel for the number of GPU assisted VMs that will be required. We are evaluating use cases at this time and have not yet obtained a clear picture of how the GPUs are being used on the pilot system.	11-Feb-21
44	Is ticking “we will comply” to all questions and putting a price enough.	Appendix A indicates that in some cases a ‘Will Comply’ response is not sufficient. In addition, fuller responses are your opportunity to lift your position in the rankings. A minimal response may very well result in you not making the pre-qualified list.	11-Feb-21

ECMWF/ItPQ/2020/307 Questions and Answers

45	<p>Please can you answer the following clarification question? 'Would ECMWF consider using the current release of Red Hat OpenStack? Red Hat OpenStack Platform (RHOSP) 16 is based on OpenStack Train API with Ussuri and Victoria cherry-picked backports to come in later releases of RHOSP 16. With RHOSP 16 Red Hat decided to only deliver LTS releases in order to better match customer requirements in terms of upgrade and Day 2 operation. Recurrent feedback to Red Hat was that upgrade procedures were too frequent and Red Hat found that end customers would not update their OpenStack environments to the latest major version of Red Hat OpenStack due to the time taken to upgrade. So instead of rebasing the customer onto a new baseline API every OpenStack release, Red Hat sticks to one baseline API (a Long term Support release) and then cherry-picks features from the two releases in between baseline API changes. RHOSP Lifecycle - <a href="https://access.redhat.com/support/policy/updates/openstack/platform">https://access.redhat.com/support/policy/updates/openstack/platform</a> RHOSP 16.x will be supported till 2024 with ELS (Extended life cycle support) extending this to 2025. Red Hat's subscription model allows its clients to use any accessible (and future) releases of its "subscribed to" offerings, until either a version reaches the end of its EOL period or the end of an ELS period. After researching both Morpheus and Terraform, they work with Train based API OpenStack deployments. To summarize the information above, if there is a need to conform completely with MHLR#1, it would make Red Hat's current OpenStack release non-compliant because the baseline API does not include Ussuri.'</p>	Please see Q42 above.	11-Feb-21
46	<p>Please can you clarify/elaborate on section 1.7.3 where it states ECMWF has a strong preference for not using proprietary licensed software or software that otherwise carries a fee for use, whether one time or subscription based - The costs of managing open source community project software is almost always higher than buying the enterprise subscription.</p>	As stated, it is our 'strong preference' not a mandatory requirement. You are advised to respond with your preferred solution and to price it accordingly.	11-Feb-21
47	<p>In Section 1.1 ECMWF state that "price will carry a significant weight during any subsequent ItT evaluation." Can you elaborate why this is the case? Is there a team that can manage a community or Open Source flavour?</p>	ECMWF is working to a budget that is pre-agreed with its Member States.	11-Feb-21
48	<p>Please can you clarify on section 1.4 - Are quarters are aligned to calendar year, fiscal year, or another. Publication of Invitation to Tender says Q2, Negotiation in Q3 and Award of contract in Q4 - Are you able to clarify on a date for each of these? This document seems to contradict these quarters - <a href="https://procontract.due-north.com/Advert?advertId=ce7d3ac8-bb5f-eb11-8106-005056b64545&amp;p=7344c2cd-e625-e711-80df-005056b64545">https://procontract.due-north.com/Advert?advertId=ce7d3ac8-bb5f-eb11-8106-005056b64545&amp;p=7344c2cd-e625-e711-80df-005056b64545</a></p>	Quarters are aligned to the calendar year i.e. Q2 is April to June, etc. ECMWF is unable to give any further refinement of project timing at this stage.	11-Feb-21
49	<p>Can you provide any more information at this stage on the Openstack that is already in pilot?</p>	ECMWF wishes to explore the widest possible range of available <b>on-premises</b> solutions. For this reason, it will not be publishing any further information about the pilot than is currently available on the portal or other online sources.	11-Feb-21

ECMWF/ItpQ/2020/307 Questions and Answers

50	3.3 'The in-house team that built the pilot consists of two cloud engineers (~ 1.5 FTE), with significant other resources made available on an as-needed basis (~ 2 FTE). The cloud engineers are RedHat RHOSP trained up to advanced level.' Can you tell us what made this pilot was successful, what if any were the challenges that were faced, what were the technical challenges, what was the end goal and how many man hours were spent to complete the POC?	Please see question 49	11-Feb-21
51	What are ECMWF plans for disaster recovery and fail over?	The Bologna Data Centre is made up of two data halls that are able to work as a redundant pair. Our plans at this time for disaster recovery do not extend beyond these two data halls.	11-Feb-21
52	What would ECMWF recovery time objective be?	Our recovery requirements at this stage are as stated in section 4.3.4.	11-Feb-21
53	Could ECMWF please clarify their thinking on the approach to pricing in the RFI response as there are a few points which seem to have conflicting intentions:1.1 Responders will, however, be required to provide indicative pricing as called for under various sections of this document. They will also be required to give an overall indicative price. (Seems to contradict 6.1 below.)1.7.3 This should include an indication of an implementation budget with breakdown for each of the features presented and an annual budget for all commercial licences. (Implementation budget would scale with the system scale so this would conflict with 6.1 below.)	ECMWF recommends that you price for scales of 100, 250 and 500 servers and that you give an annualised indicative budget over a lifecycle of 3 to 5 years.	11-Feb-21