

**ECMWF/ItpQ/2020/307**

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

**Clarifications V2**

**Issued: 28 January 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</p> <p>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</p> <p>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</p> <p>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</p> <p>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