



Invitation to Tender

Destination Earth Programme

On-demand Extremes Digital Twin

Volume II

ITT Ref: DE_330
ISSUED BY: ECMWF Administration Department Procurement Section
Date: 28 March 2022
Version: Final

Table of Contents

1	Introduction.....	3
2	Context	3
3	Contract summary	4
4	Technical specification.....	4
4.1	Data streams.....	6
4.2	Modelling systems.....	6
4.3	Prediction system set-up.....	6
4.4	Prediction system resolution.....	7
4.5	Ensemble size	7
4.6	Extremes use cases.....	7
4.7	Quality and uncertainty.....	8
4.8	High-performance computing	8
4.9	Additional requirements for software.....	9
4.10	Additional requirements for data.....	10
4.10.1	Data transfer and data formats.....	10
4.10.2	Quality control, support and documentation.....	10
4.10.3	Data and IPR	11
5	General requirements	11
5.1	Implementation schedule.....	11
5.2	Meetings.....	11
5.3	Deliverables and milestones.....	12
5.3.1	Documents and reports.....	12
5.3.2	User support	12
5.3.3	Other related DestinE activities.....	12
6	Tender format and content	13
6.1	Page limits	13
6.2	Specific additional instructions for the Tenderer’s response.....	13
6.2.1	Executive summary.....	13
6.2.2	Track record.....	13
6.2.3	Quality of resources to be deployed	13
6.2.4	Technical solution proposed.....	14
6.2.5	Management and implementation plan	14
6.2.6	Key performance indicators	15

1 Introduction

Destination Earth (DestinE) will deploy several high resolution thematic digital replicas of the Earth system to monitor and simulate natural and human activities as well as their interactions, to develop and test scenarios that would enable more sustainable developments and support European policy making. DestinE is intended to unlock the potential of observations and both physics-based and data-driven models for achieving a break-through in the realism of the simulation of Earth-system components.

The On-demand Extremes Digital Twin will develop an enhanced, configurable simulation system, informed by observations and based on a new generation of Earth system models. This enhanced system will not only allow to more realistically represent the Earth system but will also produce information at precisely those scales where the impacts of extremes are felt and where key processes are observed thus allowing users from impact-sectors to access and exploit such information for their specific application. This capability will be demonstrated targeting the near real-time quantification of risks of environmental extremes and supporting tailored, application specific adaptation and mitigation strategies. Ultimately, the digital twin will help implementing the green transition and the implementation of the European Green Deal.

2 Context

DestinE is funded by the European Commission's Digital Europe programme and implemented through a partnership between ESA, ECMWF and EUMETSAT. The present ITT only relates to ECMWF's contribution to DestinE and is further detailed in Sections 3-6.

The first phase of DestinE, the implementation phase 1, covers the period 15 December 2021 – 14 June 2024. In this phase, the main building blocks of the required infrastructure to reach DestinE's ambitious goals will be configured and deployed and their capability be demonstrated:

- The Core Service Platform (DESP; responsibility ESA) for providing a large number of users with access to observations, simulations and models, evidence-based policy and decision-making tools, applications and services, based on an open, flexible, scalable and evolvable secure cloud-based architecture.
- The Data Lake (DEDL; responsibility EUMETSAT) for handling the storage and access requirements for any input and output DestinE data that is offered to DestinE users via seamless access through the DESP including near-data processing to maximize throughput and service scalability.
- The Digital Twin Engine (DTE; responsibility ECMWF) consisting of generic software infrastructures for workflows, extreme-scale simulation and data fusion, data handling and machine learning that allow exploiting the latest digital infrastructure technology for operating Earth-system digital twins and their integration in the wider digital environment.
- The two high-priority Digital Twins (DTs; responsibility ECMWF) for generating high-quality simulations and combining simulations and observations of the Earth system at unprecedented resolution to serve the EU's Green Deal policy priorities:
 - Weather-induced and geophysical extremes DT for providing capabilities for the assessment and prediction of environmental extremes at very high spatial resolution and close to real-time decision-making support at continental, country, coastline, catchment and city scales in response to meteorological, hydrological and air quality extremes.
 - Climate change adaptation DT for providing capabilities to support climate change adaptation policy and scenario testing at multi-decadal timescales aiming at a real breakthrough in resolution at regional and national levels.

In the following phases of DestinE, these building blocks will further evolve to enhance capabilities, add thematic foci, ingest the latest scientific developments and observational information, and make use of the emerging digital infrastructure ecosystem supported by the Digital Europe programme in Europe.

3 Contract summary

This ITT covers the provision of the On-demand Extremes DT delivering a configurable, demonstration-ready capability for a flexible European monitoring and prediction framework providing decision-making support at continental, country, coastline, catchment and city scales in response to meteorological, hydrological and air quality extremes. This capability aims to combine weather, hydrology and air-quality observation and simulation capabilities at a qualitatively new level of on-demand information discovery, also providing uncertainty quantification. The demonstration of co-design with selected use-cases showing societal adaptation and mitigation options for dealing with the impacts of extreme events is included.

The On-demand Extremes DT should be technically deployed on the pre-exascale high-performance computing (HPC) facilities provided by EuroHPC. This requires, as part of this contract, to prepare for the EuroHPC computing and data handling capabilities and to demonstrate the use of these capabilities. The ITT also includes the development and implementation of software infrastructures to optimize data flow and data management, HPC performance at scale, and to demonstrate portability across HPC systems in collaboration with the DTE developments at ECMWF.

The demonstration of production-type, qualitatively enhanced extremes prediction capabilities with sufficient throughput and DestinE specific user benefits constitute the core of the requirement. Additional contributions towards enhancing both delivery performance and quality of future upgrades of the system towards operational production are also requested.

The complex nature of the activity and the dependencies on the technical infrastructure – hardware and software – may require making adjustments to development plans necessary during the period of the contract. ECMWF proposes the Agile methodology reported in Annex 4 of the Agreement template included in Volume V of this ITT, but welcomes suggestions from the Tenderer on what methodology they propose to apply for the different phases of the project based on their knowledge and experience with projects of a similar nature. The contract structure is expected to allow for change of priorities or purpose of developments, which would be addressed by this methodology. The methodologies proposed by the Tenderer must ensure that final deliverables are fit for purpose, aligned with the project vision and remain within project cost and schedule.

A single contract with one main contractor and potentially several sub-contractors towards the delivery of the On-demand Extremes DT providing the above-described capabilities is foreseen. DestinE will take ownership of the data products delivered by this contract, create further data and graphical products and make them available to the public without charge. The Tenderers are required to contribute to the generic DTE development led by ECMWF, e.g. supporting and interfacing with distributed HPC and data handling infrastructures, using common data formats and access methodologies, and contributing to a distributed software development, benchmarking, maintenance and verification framework.

Tenderers should submit proposals that demonstrate their ability to deliver high-quality, production-type capabilities with a strong foundation in advanced science, impact science, research and digital technology development. The offers will be evaluated on their scientific and technology innovation, impact-sector (e.g. energy, water, food, health) use-case relevance and progress over existing capabilities, and documented expertise and ability to fulfil production-type requirements for providing support for an effective risk management, adaptation and mitigation policy framework. The technical requirements are described in the following sections.

4 Technical specification

The successful Tenderer selected under this procurement will have a proven track-record in providing state-of-the-art predictions based on numerical weather forecasting systems, connecting such systems in configurable workflows to, for example, agricultural production, hydrology, renewable energy and air quality

sub-system models, and the ability to develop and manage complex software infrastructures. The systems are required to fully comply with the DestinE needs as defined in this section.

The required contribution consists of the following components for which further detail is provided in the following sub-sections:

1. flexible and scalable workflows for the monitoring and short-range prediction of extremes, that are configurable and operable on demand for maximizing timeliness, efficiency and benefit, for example in terms of:
 - a. geographical, event and extremes-type focus,
 - b. spatial resolution and domain size,
 - c. temporal refresh and coverage,
 - d. choice of data assimilation, ensemble size, model components and boundary conditions (the latter only for limited-area elements),
 - e. selection of observational data for monitoring, verification and (where appropriate) assimilation, ready for exploiting future data types;
 - f. inclusion of extremes-type and/or event dependent observations and simulations (e.g. hydrology models);
2. generic, observation based assessment framework for the purpose of quality assessment and for supporting uncertainty quantification;
3. demonstrated high-performance computing and data handling resource management to facilitate (1)-(2) with throughput and reliability required for production-type demonstration, e.g. considering 100s of Terabytes daily data production rates, rolling data analysis windows and different levels of storage (short/medium term), reducing and extracting targeted information and intelligent data reduction, providing detailed information on anticipated production rates (simulation days per day and Terabytes per simulated day), and contributing to data governance (portfolio documentation, integrity and provenance);
4. demonstration of value for use cases from impact-sectors such as food, water, renewable energy or air quality with a thematic/regional/national focus based on output from (1)-(2) including uncertainty quantification;
5. interfacing with ECMWF DTE developments and with DEDL and DESP as required; however, the Tenderer is not expected as part of this contract to contribute directly to the DEDL or DESP platform development;
6. provision of inputs to data and software documentation and to DestinE evaluation and quality assurance material and activities, in response to specific requests from ECMWF;
7. timely response to requests for provision of predictive modelling expertise to DestinE partners, as required by ECMWF;
8. compliance with KPIs (see Section 6.2.6) and milestones against which the progress and quality of the tendered work will be monitored;
9. upon approach by ECMWF and in agreement with ECMWF, contribution to DestinE communication and outreach activities.

In addition, Tenderers are invited to propose activities and their specific role in the future improvement of the DestinE capability provision after phase 1:

1. future ingestion of scientific and technical upgrades towards improvements in output quality and efficiency of product delivery;
2. developments supporting new products, advanced ensemble generation and uncertainty quantification towards future improvements in output quality relevant for application sectors;
3. development of a roadmap for technical upgrades of data- and workflows, code portability and performance towards a physically distributed HPC and data environment in Europe;

4. contributions to developing novel products and model components in support of risk management, adaptation and mitigation measures, together with selected users;
5. innovative contributions to material for DestinE communication, outreach and education.

Foundational scientific developments are out of scope in this activity, as are developments of methodology without prospect of production-type demonstration within phase 2 of DestinE.

The detailed definition of activities will be agreed at the contract negotiation stage with the preferred Tenderer and may be revised during the contract term. To allow efficient negotiations and minimise the complexity of any future adjustments, each development activity should be formulated, as far as possible, as a self-contained module detailed in its own work package including deliverables and milestones and by clearly assigning responsibilities.

The following sections describe specific requirements from several perspectives: technical (in terms of what components shall be delivered), capability demonstration and quality assessment.

4.1 Data streams

The On-demand Extremes DT specific data portfolio offered to users shall be developed in agreement with the DestinE partners, ECMWF, ESA and EUMETSAT, and shall also include the requirements of the use cases proposed by the Tenderer. The Tenderer shall define the complete set of variables that is required as input and output of the DT and that needs to be interfaced with the DEDL and the DESP. The data portfolio proposal shall follow the Copernicus¹ template with special consideration for the use cases, data volumes and overall data lifecycle management as proposed in this ITT, e.g. data governance, data sampling, refresh rates, compression and labelling or semantic annotation options. The data shall include uncertainty estimates towards defining application specific quality information.

4.2 Modelling systems

The prediction models (e.g. atmosphere, waves, land, selected hydrology and atmospheric composition components, options for advection of tracers in the atmosphere) shall be selected such that sufficient on-demand capabilities and configurability options can be demonstrated (compliant with Sections 4.3, 4.4, 4.5) for the benefit of prediction but also the chosen risk management, adaptation and mitigation use cases for impact-sectors (see Section 4.6) including uncertainty quantification (see Section 4.7).

4.3 Prediction system set-up

The prediction system set-up shall be compliant with or be a derivative of workflows presently in production, already delivering state-of-the art weather and environmental predictions in terms of quality, reliability and timeliness. The system shall enhance skill over present-day systems by added configurability options for the on-demand specification of, for example, geographical coverage, ensemble size, resolution and/or refresh cycles and from integrating impact-sector specific risk assessment/modelling elements (see 4.6) and include a robust approach to uncertainty quantification (see 4.7).

Monitoring and initial-condition generation capabilities shall be based on data assimilation methodologies producing an optimum estimate of the Earth system through the fusion of observations and simulations. It is encouraged to include application specific metrics in these monitoring capabilities. The Tenderer shall provide a description of which developments will be made for creating/ingesting boundary and initial conditions for the forecasts performed in the On-demand Extremes DT, and which external, non-DestinE specific components (initial and lateral boundary conditions) will be used. The Tenderer is strongly encouraged to consider a framework that utilises boundary conditions produced by the global extremes

¹ e.g. <https://cds.climate.copernicus.eu/portfolio/dataset/derived-near-surface-meteorological-variables>

digital twin developed by ECMWF in future phases of DestinE. Any requirements for input from external monitoring and prediction system output shall be described in detail and should be discussed and agreed with ECMWF at contract negotiation stage.

The observing system to be used for monitoring and assimilation shall be described with a view of upgrading the system to include novel types of observations that promise benefits for monitoring accuracy and quantifying the predictive skill of the system, but also extend into, for example, hydrology, renewable energy and air quality applications. The technical implementation shall be able to accommodate and distinguish types of observations actively assimilated and those used for monitoring and verification including new types.

Depending on the chosen impact-sector use cases (see Section 4.6), both data assimilation/monitoring and prediction components may benefit from directly integrating use-case specific impact-sector models/observations related to, for example, water run-off, river discharge, irrigation, energy yield or the dispersion of atmospheric constituents. Such integration, also through on-demand configurability shall be foreseen in the technical work.

Depending on the available HPC resources (see Section 4.8), the proposed set-up can be adapted in reduced form either by limiting the number of configurability options, the length of the predictions, by reducing spatial resolution, or ensemble size and the prediction refresh rate. It is expected that near the end of phase 1 of DestinE, sample data from predictions at target model resolution and full application diversity shall be delivered such that the progress beyond present-day simulations can be demonstrated with appropriate metrics, while the production of the entire set-up of on-demand configurability options can be staged across phase 1 and phase 2 of DestinE.

4.4 Prediction system resolution

Grid resolution is defined as the average horizontal discrete nodal point (or centre of mesh element) distance. The chosen grid resolution of the atmosphere model should be as high as possible but not coarser than 1 km, and with an appropriate choice of vertical resolution. However, coarser horizontal resolution model runs are acceptable for baseline demonstration purposes as long as the higher resolution targets to satisfy impact sector application needs can be demonstrated. A detailed roadmap shall be specified for completing the prediction system set-up with enhanced quality compared to state-of-the-art systems based on quality indicators to be defined (see also Section 4.7). The choices shall be derived from a trade-off between domain-size, spatial resolution, ensemble size (Section 4.5) and use case benefit (Section 4.6).

Spatial resolution for monitoring and quality assessment through simulation-observation data fusion shall be adjusted to the needs of the prediction system, including the demonstration of benefit beyond state-of-the-art capabilities and the potential offered by new types of observations, also available from impact sectors.

4.5 Ensemble size

If the Tenderer includes ensembles (see also Section 4.7), the target ensemble size for the prediction system is at least 20 members but ideally larger to properly represent the tails of distributions where extremes are sampled. If the Tenderer presents options for other ensemble sizes, an assessment of the impact of these options on ensemble reliability, uncertainty quantification and use cases shall be made. The ensemble configuration size will be agreed with the Tenderer following this assessment once exact HPC resource provisions have been finalised. Except for the demonstration of uncertainty quantification and the benefit for use cases, the extensive production of ensembles is not required in phase 1 of DestinE.

4.6 Extremes use cases

Selected use cases, for at least two impact-sectors such as water, food, energy resource and health management, are to be included to demonstrate the benefit of the advanced DT framework for extremes

risk management, adaptation and mitigation decision making by key stakeholders operating in these sectors at regional or local scale. For this purpose, the DT framework shall be co-designed with users from these sectors having experience with complex information systems so that tailored data and models towards a new quality of sector specific information can be demonstrated.

The use cases shall include the following capabilities highlighting the benefit from enhanced data quality and co-designing workflows:

- integration of impact sector specific models and data making use of the full space-time sampling of Earth-system states for deriving sector specific information where needed;
- sensitivity analysis of impact-sector relevant information (e.g. renewable energy production under changing wind regimes) or events (e.g. extreme events which can potentially damage renewable energy networks);
- demonstration of uncertainty quantification at use case level;
- approaches for synthesizing user oriented information on extremes facilitating decision making;
- options for and demonstration of tailored workflows for impact-sectors and selected users;
- proposals for actionable response scenarios using the On-demand Extremes DT.

The data portfolio (see Section 4.1) as well as the approach to quality and uncertainty quantification (see Section 4.7) shall be defined in close relation to the impact-sector specific requirements.

The use case assessment shall include an estimation of the specific benefits of the DT over existing capabilities.

4.7 Quality and uncertainty

The On-demand Extremes DT output shall include an objective quantification of uncertainty in terms of statistical quantities as well as a proposal for a methodology to translate this into simplified quality indicators. Ideally, the uncertainty quantification is based on well-established ensemble methods (see Section 4.5) but can also be provided by other methods such as machine learning. The quality assessment framework shall also include metrics that are specific to the use cases proposed by the Tenderer (see Section 4.6) so that the reliability of products as well as the traceability of uncertainty across the entire DT workflow can be demonstrated.

The enhanced quality of the prediction system shall be demonstrated based on available reference datasets that are considered state of the art. The choice of reference datasets shall be explained in detail. The usage of methods for exploiting the rich observational data used in numerical weather prediction models for assessing predictive skill shall be considered. The quality assessment framework shall include metrics that are specific to the use cases proposed by the Tenderer (see Section 4.6) so that both enhanced quality of products as well as the traceability of uncertainty across the entire DT workflow can be demonstrated. The quality metric framework shall be further developed throughout the project in collaboration with ECMWF, Copernicus and national services.

4.8 High-performance computing

DestinE relies on a partnership with the EuroHPC Joint Undertaking (JU). The EuroHPC JU allows the European Union and the EuroHPC JU participating countries to coordinate their efforts and pool their resources for reaching new levels of supercomputing in Europe. DestinE relies on the computing centres hosting these HPC systems (e.g. CSC, CINECA, BSC) as well as on computing centres that are organised through the partnership for advanced computing in Europe (PRACE) to create a sustainable, distributed yet well-connected European computing, data production and data exchange framework. This is complemented by the emerging DestinE DEDL and DESP infrastructure once available.

The main allocations for the HPC resources necessary for the delivery of the contract will be provided by the EuroHPC JU but can be complemented with in-kind and costed contributions from the Tenderer. Costs associated with such resources shall be fully detailed in the response by the Tenderer. The present commitment by the EuroHPC JU for activities of strategic importance for Europe (e.g. DestinE and similar activities) is 10% of the total node-hour budget available to the JU.

The Tenderer shall provide a detailed allocation estimate for both computing and data handling at the EuroHPC site of their choice. Data production and dissemination should include a separate cost estimate for ancillary HPC or cloud computing needs, and shall consider the data portfolio volume (to be agreed with DestinE partners, see Section 4.1) and data governance requirements (availability, usability, short/medium storage and persistence, etc.). The compute cost estimate shall be based on representative benchmarks that include a realistic estimate of the performance of the DT workflows, ideally on heterogeneous processor and memory hierarchy technologies as these are the primary choice for EuroHPC.

The allocation resource estimate proposal shall largely comply with criteria defined by the PRACE protocol² and:

- Demonstrate the need to use EuroHPC resources and define a platform preference (the DT codes must be suited for the requested system and are subject to a technical assessment; simulations may be redirected to a more appropriate system).
- Describe software requirements for the candidate DT workflow and provide information if these requirements are known/expected to be satisfied by the selected EuroHPC platforms. The software support required to run the DT workflows must be available on the proposed system and/or, in case of codes developed by the Tenderer these must be demonstrated to be suitable, run efficiently, and scale on selected such systems based on extrapolations from existing experience using published information of these platforms and/or based on data from actual simulations on similar HPC platforms.
- Include a proof of strong and weak scaling performance of the proposed DT workflows, ideally on platforms as close as possible to the selected EuroHPC target platforms.

4.9 Additional requirements for software

Given the challenging scale of the DestinE computing and data handling tasks, the Tenderer shall provide information on relevant software components, development priorities and a roadmap for implementing DT workflows that achieve performance and are portable across novel processor, memory and storage technologies. As much as possible, such software infrastructure implementation shall be coordinated with ECMWF as part of the DTE development that aims to implement and sustain a generic framework that will serve present and future DT developments and production. Notwithstanding the aforementioned, the Tenderer should, where possible, rely on open-source solutions.

ECMWF leads the development of the DTE functionality, which will include support for software management, continuous integration and code deployment, and benchmarking for portability and performance assessment. The Tenderer is expected to actively contribute to the software development and integrate suitable components into this framework. The DTE will also provide access to a high-performance data production and data access API, supporting data governance tasks and provide machine learning toolkits in support of data processing, data compression and computing acceleration. The Tenderer shall interface with, add to and co-develop such functionalities as relevant for the added DT and actively contribute to software management of source code contributions to shared software (e.g. git branch merging, feature developments relevant for the DT, etc.).

² <https://prace-ri.eu/wp-content/uploads/PRACE-Guide-for-Applicants-to-Tier-0-V2016.pdf>

4.10 Additional requirements for data

4.10.1 Data transfer and data formats

There are different levels of data production and data access anticipated in DestinE. The preferred solution is for data providers (including the Tenderer) to directly make use of an abstraction layer for DT data production and data streaming that is separate from the user data access provided by the DEDL or DESP. The implementation of this abstraction is the responsibility of ECMWF and it will be provided to the On-demand Extremes DT as an API to the underlying high-performance data handling service. This service will also make an interface available to the DEDL and the DESP for the subsequent dissemination and user data access process. This is part of the DTE development and requires the Tenderer to interface to these with the proposed DT models. To use this high-performance data handling service, the preferred DestinE internal data format for production is the WMO GRIB edition 2 format (GRIB2). This high-performance data access will include options for data cube³ access, extraction of partial data and support for rapid visualisation.

Beyond the GRIB2 data produced by the DestinE high-performance data handling service, additional (e.g. derived) datasets may be required. Subject to the agreement with ECMWF, these data files may be encoded in netCDF4-classic to facilitate the process of integrating datasets from other repositories. However, these data files shall at least satisfy the requirements defined by the Copernicus Data Store⁴.

In addition, the Tenderer is expected to provide a mapping capability for metadata from GRIB2 to netCDF and vice versa for all produced parameters to facilitate further processing within the DTE and to maximise the use of the On-demand Extremes DT data in other contexts. This mapping will be used also for cataloguing and for providing tools to users that transform the data between formats.

The DEDL data access model for stored datasets may require the ability for data suppliers to upload data to the DEDL, e.g. similar to ECMWF's ECPDS system, via FTP or suitable alternatives. In order to ensure that the operational service deadlines and data governance requirements are met, additional support is needed from the data providers.

The Tenderer shall consider to allocate approximately 12 person-month and define a separate work package to deliver datasets in this form as required, produce documentation of the data provenance, provide input to data catalogues, provide the metadata mapping as required, and to produce a 'manifest' of all files produced over a given period if so requested.

This approach is building on and extending the practices established and evolving during the current phase of the Copernicus Climate Data Store (CDS)⁵ service.

The Tenderer shall propose an initial data portfolio (incl. metadata) description, which can be provided as a link to a public website, as part of the On-demand Extremes DT proposal. This will be further refined during the contract phase and the Tenderer shall make an explicit provision for such action in the proposal.

4.10.2 Quality control, support and documentation

Quality control procedures (including automatic procedures) shall be established that allow to check the quality of data before transmission to the DEDL. The precise methods should be proposed by the Tenderer and will be agreed as part of the negotiations. As a minimum, the software development should follow procedures according to industry standards, e.g. provisions for version control, coherent and uniform code styles, code reviews, issue and bug tracking, branching and merge strategies, continuous unit and acceptance testing followed by continuous integration. On the data, the WMO standards for international data exchange defined in the WMO's "Manual on Codes"⁶ and set out in the WMO "Manual on the Global

³ https://en.wikipedia.org/wiki/Data_cube

⁴ <https://confluence.ecmwf.int/display/COPSRV/Dataset+delivery%3A+CDS+and+Toolbox>

⁵ <https://cds.climate.copernicus.eu/#!/home>

⁶ https://library.wmo.int/doc_num.php?explnum_id=10722

Telecommunication System” (GTS)⁷ and the standards for published metadata set out for the WMO information system (WIS)⁸ are expected.

The Tenderer is asked to document the scope of security and information management aspects to be provided and on the assets to be protected, according to ISO 27001.

In the case of ECMWF detecting possible problems with the data, providers are expected to give timely support to resolve problems quickly. Details of response procedure and time shall be established during contract negotiation.

Each DT workflow version needs to be documented at a level which defines how the data were produced to allow users to understand configurations and version changes. The datasets themselves will also need to be documented to allow later discovery by users via the DESP. The data providers will be responsible for making this documentation available as required.

4.10.3 Data and IPR

It is a condition of EU funding for DestinE that ownership of any deliverables (as defined in Volume V Agreement) developed with DestinE funding passes from the suppliers to the European Union via ECMWF. Ownership will pass from the date of creation of the result.

All pre-existing materials (e.g. software and products) used by the successful Tenderer to produce the DestinE Results will remain the property of the successful Tenderer. The successful Tenderer will have to provide a royalty-free, non-exclusive, worldwide, perpetual and irrevocable license to those pre-existing materials to the EU via ECMWF.

Developments or modifications to pre-existing materials which constitute results and are acquired or created specifically for DestinE purposes will be owned by the European Union. The pre-existing materials upon which these results are based will be licensed to the Union unless there is an agreement on the transfer of ownership to the European Union.

Upon request, suppliers may be granted a non-exclusive licence, at the discretion of ECMWF and subject to the approval by the European Commission, to use the results which they have provided to DestinE.

5 General requirements

5.1 Implementation schedule

ECMWF intends to award a single contract for a maximum duration of 22 months, expected to commence by mid-July 2022, and with an end date not later than 30 April 2024.

The successful Tenderer is expected to provide a detailed schedule as part of the tender response. The proposed time plan and schedule shall address the main tasks, inputs, outputs, intermediate review steps, milestones and deliverables. A roadmap of future developments beyond the contracted period is also highly desirable.

5.2 Meetings

Regular progress meetings will be held with ECMWF during the contract to assess contract status, risks and actions. ECMWF will organise annual meetings to bring together all DestinE capability providers, of which the first will be at M12 after kick-off. The successful Tenderer is expected to attend these meetings. The successful Tenderer is also expected to attend monthly (video-conferencing) meetings to discuss progress

⁷ https://library.wmo.int/doc_num.php?explnum_id=10469

⁸ <https://community.wmo.int/activity-areas/wis/wis-manuals>

and other topics that cut across different aspects of DestinE. The cost of attending these meetings shall be covered by each successful Tenderer and shall be included in the tendered price. ECMWF may adjust meeting frequency as needed.

In addition, the successful Tenderer is expected to participate in a technical working group – which may also include other DestinE partners and relevant collaborators – aimed at discussing issues related to product definition, generation and integration of the overall DestinE infrastructure. These discussions will be convened at regular intervals through video-conferencing.

5.3 Deliverables and milestones

Deliverables are to be defined by the Tenderer based on the requirements outlined above. They can be in the form of documents or reports, datasets and support to users or other related DestinE activities. Also requirements related to the delivery of software and data have been described above (see Sections 4.9 and 4.10). The requirements for all other types are described in the following subsections. The Tenderer is encouraged to limit the number of deliverables and milestones to a reasonable amount so that their preparation, review and revision remains manageable.

Each deliverable shall have an associated resource allocation (person-months and financial budget). The total of these allocated resources shall amount to the requested budget associated with payroll as detailed in Volume IIIA of this ITT.

Milestones should be designed as markers of demonstrable progress in capability development and/or quality of capability delivery, as applicable. They should not duplicate deliverables.

5.3.1 Documents and reports

All project reports shall be produced in English. The quality of reports and deliverables shall be equivalent to the standard of peer-reviewed publications and practice. Unless otherwise specified in the specific contract, deliverables shall be made available to ECMWF in electronic format (Microsoft Word/PDF/Microsoft Excel or compatible), via the DestinE Deliverables Repository portal; the details will be agreed at the negotiation stage.

Please refer to Clause 2.3 and the Annex 5 of the Volume V Agreement for details on Reporting Obligations.

5.3.2 User support

It is expected that EuroHPC will provide dedicated user support for use of their facilities. ECMWF will lead the DTE development, provide APIs and respond to bug reports, and provide limited support for the development and interfacing tasks with the DTs as well as individualised support on technical queries related to data formats and data access.

In addition, ECMWF will make available DestinE staff resources for providing technical and scientific support to address questions related to DT development, interfaces to data workflows, machine learning toolkits, monitoring facilities and the interface to DEDL and DESP developments.

The Tenderer is expected to contribute to the delivery of technical support for the data and/or capabilities they provide. Such technical user support shall take the form of a direct response to individual user queries as required, as well as potential contributions to FAQs, user guides and knowledge bases. The Tenderer shall cost this as a separate task within the work package defined in Section 4.10.1.

5.3.3 Other related DestinE activities

The successful Tenderer is required to support the wider DestinE activities, for example the DestinE partnership activities, communication and training and outreach. Sufficient resources for covering these aspects shall be foreseen.

Outreach activities will be organised by ECMWF during the period of the contract. In such instances, the contractors will be approached by ECMWF for support on developing and delivering contents. Similarly, DestinE will require contributions to training material on environmental prediction, and risk management, adaptation and mitigation for the selected use cases from the contractor.

Contractors shall not establish their own brand for the selected projects but rely on and use DestinE and ECMWF branding. A communications package (including guidelines, logos and templates) will be provided by ECMWF at the start of the contract.

6 Tender format and content

General guidelines for the tender are described in Volume IIIB. Specific requirements to prepare the proposal for this particular tender are described in the next sub-sections.

6.1 Page limits

As a guideline, it is expected that individual sections of the Tenderer’s response do not exceed the page limits listed below. These are advisory limits and should be followed wherever possible, to avoid excessive or wordy responses.

<i>Section</i>	<i>Page Limit</i>
<i>Executive Summary</i>	2
<i>Track Record</i>	2 (for general) and 2 (per entity)
<i>Quality of resources to be Deployed</i>	2 (excluding Table 1 in Volume IIIB and CVs with a maximum length of 2 pages each)
<i>Technical Solution Proposed</i>	30 (Table 2 in Volume IIIB, the section on references, publications, patents and any pre-existing IPR is excluded from the page limit and has no page limit)
<i>Management and Implementation</i>	10 (excluding Table 4 and Table 5 in Volume IIIB) + 2 per each Work package description (Table 3 in Volume IIIB)
<i>Pricing Table</i>	No limitation

Table 1: Page limits

6.2 Specific additional instructions for the Tenderer’s response

The following is a guide to the minimum content expected to be included in each section, additional to the content described in the general guidelines of Volume IIIB. This is not an exhaustive description and additional information may be necessary depending on the Tenderer’s response.

6.2.1 Executive summary

The Tenderer shall provide an executive summary of the proposal, describing the objectives, team and summarising the proposed technical solution and capability demonstration.

6.2.2 Track record

The Tenderer shall demonstrate for themselves and for any proposed subcontractors that they have experience with relevant projects in the public or private sector at national or international level. ECMWF may ask for evidence of performance in the form of certificates issued or countersigned by the competent authority.

6.2.3 Quality of resources to be deployed

The Tenderer shall propose a team providing the skills required for developing and demonstrating the solutions complying with technical requirements set out in Section 4. The team shall include a dedicated Project Manager with at least 5 years of experience in management of large-scale projects. The Tenderer

shall describe the experience of the Project Manager and the technical project team in performing activities related to the various aspects of this tender.

6.2.4 Technical solution proposed

The Tenderer shall give an introduction to the overall proposed technical solution to demonstrate their understanding of the DestinE context and the specific requirements of the present tender. This section shall also include information on other third-party suppliers that are proposed for delivering the technical solution.

6.2.4.1 Existing capabilities

Tenderers should present information outlining the strength of their present capabilities in the following form:

- A brief description of the scientific heritage of the proposed technical components (such as simulation models and data fusion methods), also describing their current level of maturity.
- Peer-reviewed publications or other documentation describing the system's scientific performance.
- A brief description of experience with large HPC systems and supporting software stacks, and developing and running complex systems at scale.
- A brief description of the existing data handling system and the capacity to handle big data.

6.2.4.2 Provision of technical capabilities

Tenderers shall describe in detail their proposed technical approach to develop, implement and demonstrate the technical solution in compliance with the technical requirements laid out in Section 4. The description shall include information on how the proposed solution maps onto the requirements formulated throughout this document.

Tenderers shall describe their plan to maintain and update existing capabilities in the course of the contract. If any new, scientific developments are considered necessary during this period, these shall also be described, for context, even if they may not qualify for funding under this contract (see in- and out-of-scope topics in Section 4). Where insufficient information does not allow the fully compliant description of a technical solution, such case shall be highlighted to facilitate discussions at the negotiation stage.

6.2.4.3 Computing and data handling resources

Tenderers shall provide information about their preferred (EuroHPC) computer system for deployment of the technical solution and describe the benefits of the capabilities of this system with respect to the requirements of the contract. Specific needs with respect to HPC allocations and ancillary needs in terms of cloud computing, processing and data handling, data storage and data transmission are to be specified separately (see also Section 4.8).

If other computing resources (guaranteed in-kind, subject-to-proposal in-kind, or charged to this contract) are provided, the estimate of the associated computing and data handling cost needs to be accompanied by information on the nature of the anticipated resource and the elements included in this cost, in as much detail as possible.

6.2.5 Management and implementation plan

The Tenderer shall provide a detailed implementation plan of proposed activities for the duration of the contract. Deliverables should be consistent with the technical requirements specified in Section 4.

The Tenderer is requested to include management and implementation activities within a dedicated work package (WPO). The number of milestones is not prescribed, but they should be designed as markers of demonstrable progress in capabilities development and/or quality of capability delivery to keep progress monitoring manageable.

Adjustments to the proposed implementation plan can be proposed by the successful Tenderer once per year (except for force majeure), depending on the needs for the evolution of the technical solution, changed user requirements, or other requirements, but must be agreed to by ECMWF and between the European Commission and ECMWF.

As part of the general project management description the Tenderer shall consider the following elements (this is not an exhaustive list):

- Semestrial, annual and final reports shall be provided in accordance with the Volume V Agreement Clause 2.3 and Annex 5.
- An annual work plan is expected to be agreed at negotiation for 2023. The annual work plan for 2024 shall be provided in July of 2023.
- Monthly video-conferencing with ECMWF and a proposal for involvement of ECMWF in major project reviews shall be provided as part of the management plan. The contractor is responsible for the organisation of such meetings, including provision of minutes.
- If relevant, a list of sub-contractors and details of their contribution, key technical personnel involved in the contract, legal names and addresses shall be provided. The tenderer shall describe how the Volume V Agreement, in particular Clause 2.9, has been communicated to all their sub-contractors.
- The Tenderer shall describe in the Proposal the management of personal data and how this meets the requirements of Clause 2.8 and Annex 6 of Volume V Agreement.

The table below provides the template to be used by the Tenderer to describe the complete list of deliverables, milestones and schedules for this work package. All milestones and deliverables shall be numbered as indicated. All document deliverables shall be periodically updated and versioned as described in the table. Tenderers shall provide preliminary versions of the completed tables as part of their bid.

Deliverables for this work package shall include the following reports:

WP0 Contractual Obligations Template			
<i>#</i>	<i>Nature</i>	<i>Title</i>	<i>Due</i>
D0.y.z-YYYY	Report	Semestrial Implementation Report (January-June YYYY) YYYY being the Year n This includes a specific Financial Report	Annually on 15/07
D0.y.z-YYYY	Report	Annual Implementation Report YYYY YYYY being the Year n-1 This includes a specific Financial Report	Annually on 15/01
D0.y.z	Report	Final Implementation Report	60 days after end of contract
D0.y.z-YYYY	Report	Annual Work Plan YYYY YYYY being the Year n+1	Annually on 31/08
D0.y.z-YYYY	Other	Copy of prime contractor's general financial statements and audit report YYYY YYYY being the Year n-1	Annually (no-cost associated)

6.2.6 Key performance indicators

Contractors shall report to ECMWF on a set of Key Performance Indicators (KPIs) suitable for monitoring various aspect of service performance, including (but not limited to):

- Capability development
- Capability demonstration
- Output quality
- Technical performance
- User support

The KPIs, to be defined by the Tenderer, are subject to review by ECMWF and may be updated if necessary. The Tenderer should propose KPIs such that the overall DestinE KPIs (see below) can be reported upon where applicable.

KPI	Method	Unit	Applicable Phase
Partnerships			
Hydrology Applications	Agreement implementation	Number	I-IV
Energy Applications	Agreement implementation	Number	I-IV
Food Applications	Agreement implementation	Number	II-IV
Health Applications	Agreement implementation	Number	II-IV
Trans-continuum	Agreement implementation	Number	I-II
Others	Agreement implementation	Number	II-IV
Digital Twin capabilities			
Spatial resolution/coverage of monitoring and prediction DT Extremes datasets	Monitor system usage/evolution; define DT improvement over existing systems	km/area	I-IV
Spatial resolution/coverage of monitoring and prediction DT Climate datasets	Monitor system usage/evolution; define DT improvement over existing systems m	km/area	I-IV
Temporal availability/timeliness of DT output for continuous production mode	Monitor system usage/evolution; define DT improvement over existing system	hours	I-IV
Temporal availability/timeliness of DT output for on-demand production mode	Monitor system usage/evolution; define cost-benefit of DT set-up improvement over existing systems	hours	I-IV
Availability of decision-ready information derived from DT output	Monitor DT output uptake per impact sector	%	II-IV
Digital Twin data uptake			
Number of (service-level) applications using DT-Extremes output	Monitor the actively running applications on core platform using output (simulations, observations)	Number	II-IV
Number of (service-level) applications using DT-Climat e output	Monitor the actively running applications on core platform using output (simulations, observations)	Number	II-IV
Number of (service-level) applications using full-resolution, high-frequency output	Monitor the actively running applications on core platform using output (simulations, observations)	CPU Hours	II-IV
Number of (service-level) applications using critical-path output	Monitor the actively running applications on core platform using output (simulations, observations)	Number	II-IV
Number of applications producing candidate models to be added to DT Engine	Monitor and evaluate the number of applications suitable for integration	Number	II-IV
Number of new datasets created from DT output	Monitor the datasets made available from core platform relative to original portfolio	Number	II-IV
Computing and data handling performance			
Amount of data sets pushed into the Data Lake(s)	Monitor the overall data flow across bridges	Number, TB	I-IV
Ingestion, access & usage of Earth observation data	Monitor the overall data flow across bridges and uptake by DT Engine	Number, TB	I-IV

Actual HPC node allocation for continuous production mode	Monitor the node-hour allocation on test systems and EuroHPC platforms	Node-hours/day	I-IV
Actual HPC node allocation for on-demand product mode	Monitor the node-hour allocation on test systems and EuroHPC platforms	Node-hours/day	I-IV
Sustained vs peak performance	Assess application specific sustained performance on test systems and EuroHPC platforms	%	I-IV
Digital Twin Engine			
Extreme-scale software component uptake	Monitor number of models/data assimilation systems employing DTE modules	Number	II-IV
Machine-learning software component uptake	Monitor number of models/data assimilation systems employing DTE modules	Number	II-IV
HPC efficiency gains in DT production	Monitor the change in time-to-solution/reduction of node allocations	%	II-IV
Data exploitation gains in DT production and use in applications	Monitor the change in data touched and actively used	%	II-IV