



Subject:	Framework for Member State time-critical applications
Action:	To consider and comment on the proposal to develop a framework for Member State time-critical applications

## 1. Background

A number of Member State applications running on the Centre's computing systems require access to real-time data produced by the latest ECMWF operational analysis or forecast. A mechanism called "job submission under SMS control" (see Section 2 for more details) was developed several years ago to allow the submission of Member State jobs once the required data became available. This mechanism has usually been used to access data produced by one of the ECMWF models, possibly post-process them and then send them to the relevant Member State(s).

This mechanism has allowed the creation of tailored products that require the post-processing of large quantities of data that would otherwise have been impossible to transfer, due to the bandwidth limitations of the ECMWF wide area network. For example, "job submission under SMS control" has been used by some Member States to create tailored products based on ECMWF EPS data.

In recent years the Centre has been asked to support some Member State applications that rely on accessing specific ECMWF data for their time-critical activity (see Section 3 for more details). One example is the COSMO-LEPS project; another is the recent Greek NMS request to support their LAM forecasting activities at ECMWF during the Olympics and Para-Olympics in August-September 2004.

All such Member State projects have differing characteristics and vary in their duration (from a few months to a few years), their complexity and the effort required for their setting up. In the past, similar requests have been satisfied with ad-hoc technical solutions.

The aim of this paper is to propose a framework that would allow Member States to run time-critical applications themselves, reducing the need for individual ad-hoc solutions. The proposed framework is described in Section 5.

## 2. Current options for Member State time-critical applications

Member States interested in running regular time-critical applications on the Centre's computing systems currently have two main options. The first is to use a simple "job submission under SMS control". The second is to submit a request for an ECMWF Optional Project.

## 2.1 Job submission under SMS control

This mechanism allows Member State users to set up jobs in advance, so that they are automatically submitted, once certain points in one of the daily ECMWF operational suites have been reached. The main purpose of this method is to ensure that the latest analysis or forecast data required by a job are available, before that job is submitted. This facility runs on the Member States' general-purpose server (ecgate) and makes use of SMS, the Supervisor Monitor Scheduler, for the jobs' submission. A user provides the job and indicates at what stage of the operational suite it is to be submitted. When the operational suite has reached the specified point, for example a predefined time-step in the 12 UTC deterministic forecast, the Supervisor launches a task that submits all the relevant jobs on behalf of their file owners. The jobs are executed using the appropriate username under normal ecgate scheduling and they are not given any specific priority. The jobs' owners are fully responsible for these jobs, in particular for their development, maintenance, and monitoring. For more detailed documentation on this method, see [www.ecmwf.int/services/computing/docs/ms\\_items/MS\\_jobs\\_via\\_SMS.html](http://www.ecmwf.int/services/computing/docs/ms_items/MS_jobs_via_SMS.html).

The simple "job submission under SMS control" mechanism is used by many users from many different countries: currently SMS submits about 300 jobs belonging to approx. 60 different Member State users daily. These jobs access the output of one of the ECMWF models and then send it to the Member States either directly, using standard ftp or ectrans, or after some post-processing on one of the Centre's computing systems.

## 2.2 ECMWF Optional Projects

The procedure for requesting and establishing an ECMWF Optional Project was adopted by Council at its 30th session (May 1989) (ECMWF/C/30/M(89)1 paras. 51-61 and Annex III); and is described in detail at [www.ecmwf.int/about/basic/volume-1/procedure\\_for\\_optional\\_projects/](http://www.ecmwf.int/about/basic/volume-1/procedure_for_optional_projects/). Section 1. The Introduction to this procedure states:

*An 'ECMWF Optional Project' is a project, proposed by a Member State or group of Member States, in which all Member States participate apart, from those that formally declare themselves not interested in participating therein.*

*Before agreeing to an Optional Project, the Council will wish to satisfy itself that:*

- *the Project will contribute to the objectives of the Centre;*
- *provision has been made for any computer resources required;*
- *provision has been made for any manpower resources required, and*
- *any financing required will be made available by the participants.*

Interested Member States are required to submit a proposal which includes two parts: one (Technical & Scientific) listing the Project objectives, together with the technical and scientific content, the timetable and an estimate of the computer resources required and associated manpower resources (if any); the other part (Financial) describing the financial and budgetary aspects. Computing resources for Optional Projects are usually provided from the normal allocations of the participating Member States.

Following Council approval the Centre implements Optional Projects as separate operational suites to limit their possible impact on ECMWF core activities. Absolute priority is given to the Centre's core activities and, if problems affecting the core activity suites arise or long system sessions are needed, the optional project suite may be delayed or abandoned (and, if required by the project specification, backup products disseminated). The Centre is fully responsible for the optional projects' suites, in particular for their development, maintenance, migration, monitoring, and call-out support.

Since Council adopted the procedure for ECMWF Optional Projects, three such Projects have been approved: the Optional Project for the prediction of ocean waves (approved by Council at its 34th session (June 1991)) which then became a core ECMWF project on 1 January 1999, the Optional 00 UTC Forecast Project (also approved by

Council at its 34th session) and the Boundary Conditions (BC) for Limited Area Modelling Optional Project, which replaced the 00 UTC project and was approved by Council at its 50th session (June 1999).

### **3. Recent requirements**

In recent years the Centre has been asked to support specific Member State activities that are routinely run on the ECMWF computing systems. The COSMO-LEPS project is an example, which has been running at the Centre since November 2002. When the project was established, it required access to limited area products from all members of three consecutive runs of the ECMWF EPS. As it was not feasible to disseminate such a large volume of products, the project was implemented on the ECMWF computing system. The participating Member States created their own suite, using the Supervisor Monitor Scheduler in a similar way to the operational suite run by ECMWF. The suite is triggered by the same "job submission under SMS control" mechanism and is therefore linked to the running of the ECMWF operational suite. The suite was developed and is maintained by the COSMO-LEPS group, which also looks after its routine running. The Centre has been mainly involved in providing advice and support during the set-up phase of the project and for the development of the suite. For more information on this project, see ECMWF Newsletter no. 98, Summer 2003, available at [www.ecmwf.int/publications/newsletters/pdf/98.pdf](http://www.ecmwf.int/publications/newsletters/pdf/98.pdf).

A similar approach has been used by the UK MetOffice for the implementation of their Monthly Outlook suite, which has been running at the Centre since 2003. The suite was developed at the Centre because it requires the processing of a considerable volume of data and, furthermore, the required data were not available within the dissemination.

More recently, in May 2004 the Centre received a request from Norway to support LAMEPS, their limited area ensemble prediction system using HIRLAM. LAMEPS needs perturbations for the lateral boundaries, as well as for the initial conditions. In past years Norway has tested LAMEPS, using both the normal EPS model output and its targeted version, TEPS. Several tests established that the version of LAMEPS using TEPS verified better than the one using EPS. Norway has asked the Centre to support them in setting up on the ECMWF supercomputer a semi-operational run of a TEPS system that will be used to initialise their LAMEPS running in Norway. The plan is to run this activity twice a week. The computing resources needed to run the TEPS system would be covered by the normal Norwegian allocation. The request from Norway is currently being evaluated by the Centre. One option is to assist Norway in setting up a TEPS suite, which they would run semi-operationally. The other is to run the TEPS suite as an annex to the ECMWF operational suite. This latter option could constitute a pilot project for an enhanced service that the Centre could offer to all Member States and is described below.

Germany has also enquired about the provision of specific support for their time-critical jobs, which are run on a daily basis on the ECMWF computing facilities. These jobs use both the Member State server, eagate, to retrieve specific IFS model output, and the supercomputer to interpolate it to the GME grid. The resulting data is then archived into MARS and also sent to DWD.

Finally, the Greek NMS recently asked the Centre to support their LAM forecasting activities at ECMWF during the Olympics and Para-Olympics in August-September 2004. This project required the provision of hourly boundary conditions from the BC Optional Project runs. In this case, the Centre has, on an experimental basis, for the duration of the Olympics, provided hourly boundary conditions to all countries participating in the BC Optional Project.

### **4. Limitations of the current options**

The main limitations of the simple "job submission under SMS control" for time-critical applications are:

- the only service provided is the submission of the relevant Member State jobs; if, for any reason, a job fails, it is the responsibility of the job's owner to resubmit the job, after having fixed any problems.

This is the case, even if the job fails due to the temporary unavailability of any service (e.g. batch system down, file system full, etc.);

- the Member State jobs submitted via this mechanism cannot be differentiated from any other Member State jobs; they are not easily identifiable and, for example, ECMWF operators are uninformed of where these jobs run and what they are supposed to do;
- there is no specific support for any of these jobs: they can be affected by software and system changes implemented at ECMWF;
- in some cases these jobs rely on services for which no backup system is available (e.g. ecgate, hpca).

The main limitation of the Optional Project framework derives from the administrative effort involved in setting up the individual jobs. For this reason, any Project should normally have an expected life of at least a few years.

## **5. Proposed framework for Member State time-critical applications**

To overcome some of the limitations described in Section 4 and to offer a service more suitable for the recent Member State requests, the Centre proposes the following three options:

### **5.1 Simple job submission monitored by the Centre**

ECMWF proposes to offer an enhanced job submission facility to provide a more suitable service for simple Member State time-critical applications. The new facility will be based on the ECaccess framework and, in particular, will use the ECaccess web interface. Member State users will be able to link the submission of their jobs to the occurrence of a predefined number of events (e.g. availability of the 12 UTC 4D-Var analysis, availability of day 10 of the 00 UTC based forecast, etc) in the ECMWF operational suite. For every such job a corresponding SMS task and the relevant triggering will be automatically created. When the data required from the ECMWF operational suite are available, all the SMS tasks relevant to Member State jobs will be started. The use of SMS will allow ECMWF operators to identify and more easily monitor this kind of activity. ECMWF operators will also be able to resubmit such jobs, if they fail because of the temporary unavailability of some essential service, for example, they will be able to resubmit jobs that fail due to temporary file system problems, such as file systems becoming full or mounting problems. Member State users will be able to monitor their jobs using the standard ECaccess job monitoring interface.

Member State users who submit their jobs using this facility could be asked to provide instructions for the ECMWF operators, describing what to do in case of failure, the call-out level of any job and the contact details of a person or persons who could be contacted, if necessary. This enhanced facility would allow every registered Member State user to submit a job, once the required data, produced by the ECMWF operational suite, became available, and would benefit from some basic monitoring and restart services performed by the ECMWF operators. Member State users would still be fully responsible for these jobs, in particular for their development, maintenance, on-call support and migration.

The Centre considers that this facility should be suitable for most simple Member State users' jobs, including the time-critical jobs run by Germany (GME grid conversion). The Centre plans to have this enhanced facility available for testing by Member State users in the first half of 2005.

### **5.2 Member State SMS suites monitored by the Centre**

For more complex applications requiring several tasks with interdependencies among them, the Centre recommends that Member State users should create their own SMS suite and then trigger its execution by a simple job submitted using the enhanced job submission facility described above. This suite can then be made visible to the ECMWF operators, who can take actions as offered for the enhanced job submission facility.

The suite should be developed according to technical guidelines which will be provided by the Centre. As with the previous option, Member States who develop such suites could be asked to provide instructions for the

ECMWF operators, describing what to do in case of failure, the call-out level of any job, and the contact details of a person or persons who can be contacted in case of problems. The Centre considers that this option should be suitable for Member State applications that have a degree of complexity similar to the COSMO-LEPS project or the UK MetOffice Monthly Outlook suite.

Member States interested in using this facility would need to submit a detailed request to the Centre via their TAC representative.

### **5.3 Member State SMS suites managed by the Centre**

The Centre is investigating the possibility of running some of these complex Member State applications as separate suites under the following clearly defined conditions:

- such applications are developed, tested and maintained by the Member States according to technical guidelines which will be provided by the Centre; in particular, it must be possible to test and run the applications using ECMWF pre-operational (e-suite) data before an ECMWF e-suite is implemented operationally;
- after the application has been running regularly in general user mode, it is handed over to ECMWF who will test it as a pre-operational suite (e-suite). If problems are encountered during this phase, the application will be handed back to the Member State for them to solve the problems. This process will be repeated until the application runs stably in pre-operational mode;
- after the successful completion of pre-operational testing, the Centre will start running the application as a separate operational suite to limit any possible impact on the core ECMWF activities;
- any change to the application will, in general, require a new phase of pre-operational testing;
- Member States will provide instructions describing what to do in case of failure, the call-out level of any job and the contact details of the person providing second-level support;
- ECMWF operators will monitor the running of such Member State suites;
- ECMWF could, if required, provide first-level on-call support, while the second-level support would be provided by the Member State responsible for the application;
- Member States will be responsible for the migration of the application, if this is required, following changes in the ECMWF operational suite or in any of the ECMWF computing systems used by the application;
- the normal Member State computing resources allocation will be used to run these applications.

The Centre envisages that such an option would be suitable for Member State activities such as the Norwegian TEPS system or the COSMO-LEPS project, if these activities are to be run under ECMWF control. Member States interested in using this facility would need to submit a detailed request to the Centre via their TAC representative.

The Centre considers that this option would be suitable when one or a small number of Member States request the Centre to run a specific time-critical project. If most but not all Member States are interested, the ECMWF Optional Project mechanism will generally be more appropriate. Optional Projects typically last several years and ECMWF supports them for their duration. . However, ECMWF would reserve the right to give 12 months' notice of withdrawing its support from projects that fall within "the SMS suites managed by the Centre" category, should technical problems or resource constraints arise. The principal characteristics of the proposed options for Member State time-critical applications are summarised in Annex 1.

In general, the Centre's core operational activities will always be given higher priority and, if problems affecting the core activity suites arise or long system sessions are needed, the Member State suites and jobs may be delayed. Member States may need to consider the possibility of setting up suitable backup procedures for such eventualities.

## **6. Impact on human resources**

The service described in Section 5.1 will require some human resources for the development of the new facility and will increase the workload on the ECMWF operators who will monitor these Member State jobs. However, the Centre anticipates that this option can be implemented and supported within the already existing ECMWF staff resources.

ECMWF User Support has, in essence, already been providing the service described in 5.2 for several years. This type of workload has been steadily increasing and is part of the requirement for an additional consultant in User Support. Such an additional consultant is requested in the Four-Year Programme of Activities. (see document ECMWF/TAC/34(04)16, item 14.3).

The service described in 5.3 can be provided with the existing human resources, provided that it does not go beyond a small number of simple Member State projects. If a larger number or more complex projects need to be supported, additional staff resources would be required. Depending on the number or complexity of the applications, either 50% of a person or a full-time person would be required. The additional resources would have to be distributed across User Support, Meteorological Applications and Computer Operations.

## **7. Action required**

The Committee is invited to consider and comment on the proposal to develop a framework for Member State time-critical applications.

*Principal characteristics of the proposed options for Member State time-critical applications*

	<b>Requested by</b>	<b>Developed by</b>	<b>Tested by</b>	<b>Migrated by</b>	<b>Monitored by</b>	<b>On-call support by</b>
<b>Simple job submission monitored by the Centre</b>	does not require any specific request; any registered user with access to ECMWF computing systems and access to real-time data can use this option	the user who owns the specific job	the user who owns the specific job	the user who owns the specific job	ECMWF operators	the user who owns the specific job
<b>Member State SMS suites monitored by the Centre</b>	TAC representative of the relevant Member State	the user who is responsible for the Member State suite	the user who is responsible for the Member State suite	the user who is responsible for the Member State suite	ECMWF operators	the user who is responsible for the Member State suite
<b>Member State SMS suites managed by the Centre</b>	TAC representative of the relevant Member State	the user who is responsible for the Member State suite	tested in general user mode by the user who is responsible for the Member State suite; tested in pre-operational mode by ECMWF staff	the user who is responsible for the Member State suite	ECMWF operators	first-level support provided by ECMWF staff; second-level support provided by Member State staff
<b>ECMWF Optional Project</b>	Requested by a Member State or a group of Member States and then approved by Council	Staff resources provided as part of the Optional Project	Staff resources provided as part of the Optional project	Staff resources provided as part of the Optional project	ECMWF operators	ECMWF staff in conjunction with the staff resources provided as part of the Optional project