BARCELONA DUST FORECAST CENTER: TECHNICAL REPORT 2014

BDFC-2015-001

F. Benincasa, K. Serradell, E. Terradellas, J. M. Baldasano and S. Basart

Barcelona Dust Forecast Center (BDFC)

E-mail: dust@aemet.es

31 julio 2015

TECHNICAL REPORT

Series: Barcelona Dust Forecast Center (BDFC) Technical Report

A full list of BDFC Publications can be found on our website under:

http://dust.aemet.es/about-us/technical-reports

Barcelona Dust Forecasting Center (BDFC)

C/Jordi Girona, 29 | 08034 Barcelona (Spain)

Summary

The present document describes the technical configuration of the Barcelona Dust Forecast Center and its main vulnerabilities.

Contents

1. Introduction

In May 2013, the 65th Session of the WMO Executive Council designated the consortium formed by the State Meteorological Agency of Spain (AEMET) and the Barcelona Supercomputing Center (BSC) to create in Barcelona the first Regional Specialized Meteorological Center with activity specialization on Atmospheric Sand and Dust Forecast (RSMC-ASDF). The Center would operationally generate and distribute dust predictions for Northern Africa, Middle East and Europe.

The Barcelona Dust Forecast Center (BDFC, http://dust.aemet.es) was created in February 2014 to be this RSMC-ASDF and was publicly presented on 10 June 2014. It is located in BSC facilities. Thus, it takes advantage of the technical infrastructure and user support of its host institution. This document describes the technical infrastructure of the BDFC and complements the BDFC Activity Report 2014, available on the Center's website: http://dust.aemet.es/about-us/BDFC-2015-001.pdf.

The main components of the Center are:

- Integration of the numerical dust forecast models
- Generation and archive of final products
- Dissemination of products through the internet
- Dissemination of products through the WMO GTS/WIS

1.1. Regional Center Staff

Currently, the Regional Center staff is the following:

Enric Terradellas, technical director
Sara Basart, research and operations
Francesco Benincasa and Kim Serradell, technical support
José M. Baldasano and Emilio Cuevas, scientific advisers

1.2. Users

The BDFC performs a quarterly monitoring of the accesses to its web portal. The results (Table 1) show a steady increase in the number of users, sessions and page views.

Table 1 BDFC website: Quarterly number of users, sessions and page views

Season	Users	Sessions	Page views
Spring (Mar-May)	422	878	2,521
Winter (Jun-Aug)	1,132	2,449	6,134
Autumn (Sep-Nov)	1,909	3,175	6,901

1.3. Service disruptions

In 2014, the BDFC has experienced three kinds of disruptions:

a. Web service outages due to excessive concurrent accesses (software failure on the BDFC main server).
b. Daily forecast not received on due time due to a failure on the Mare Nostrum (3 times). It has not occurred since the model run has been duplicated.
c. Service blackout due to scheduled infrastructure maintenance.

It can be highlighted that there has been no interruptions or delays in the services between August 20 and end of year. Individual disruptions are listed in Table 2.

Table 2 Service disruptions experienced in 2014.

Date	Technical issue
19-20 Aug	Services down for 27 hours for electrical maintenance work in the campus (previously scheduled and announced) (3)
27 Apr	Delayed forecast availability (2)
27 Apr	Web service down for 9 hours for maintenance work in the campus (previously scheduled and announced) (3)
22 Apr	Delayed forecast availability (2)
19 Apr	Delayed forecast availability (2)
2-3 Apr	Web service down between 2 Apr 17:00 UTC and 3 Apr 09:00 UTC due to excessive concurrent access (1)

2. Technical configuration

Starting on 25 March 2014, the BDFC generates and distributes daily forecasts (on a 3-hourly basis up to a 72-hour lead time) of the dust contents in the atmosphere using the NMMB/BSC-Dust model. The geographical domain includes Northern Africa (north of equator), Middle East and Europe (see Figure 1).

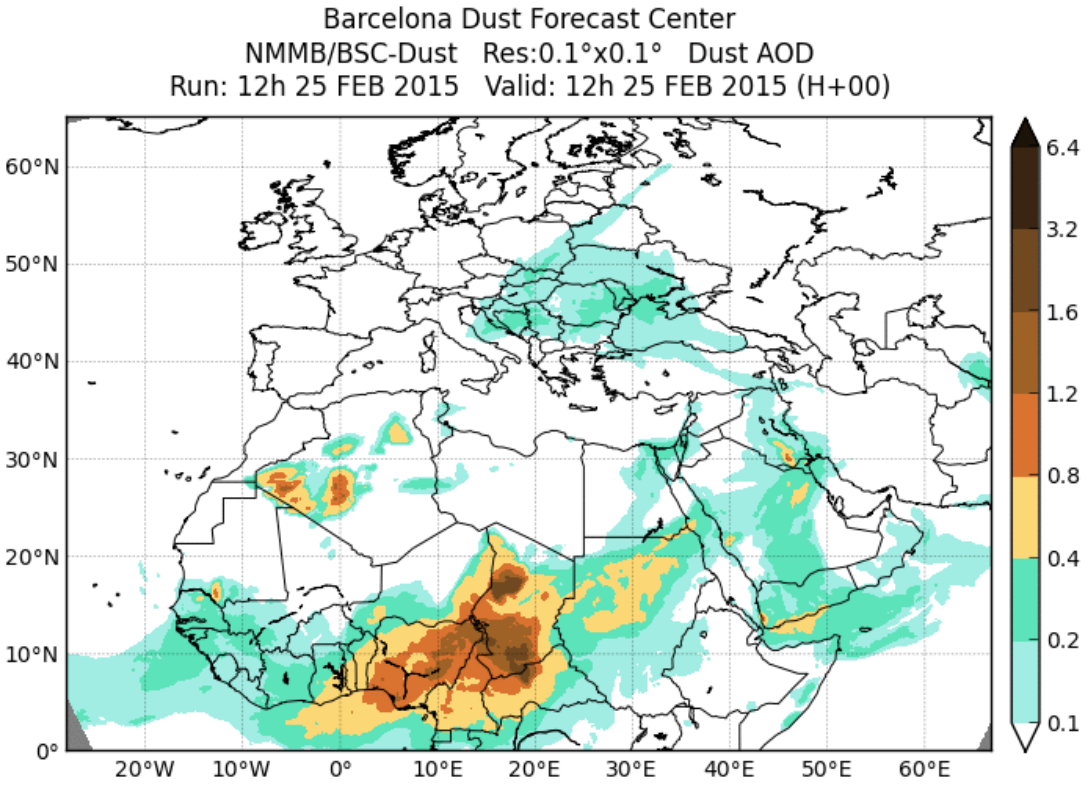

Figure 1. Domain of simulation of the NMMB/BSC-Dust model. The colour scale shows AOD (at 550nm) forecast on 25th February at 12h.

Two model integrations are daily performed:

- The 0.10° latitude x 0.10° longitude x 40 σ-layers resolution forecast (hereafter referred as 0.10°) run at the Mare Nostrum III Supercomputer using dedicated resources to ensure the daily execution. This is the default forecast.
- The 0.33° latitude x 0.33° longitude x 24 σ-layers resolution forecast (hereafter referred as 0.33°) run at a dedicated server (forecast server). This is the backup forecast.

Both model configurations use initial meteorological conditions from the U. S. National Centers for Environmental Prediction (NCEP) global analysis at a 0.5° latitude x 0.5° longitude

resolution and 6-hourly boundary meteorological conditions from the NCEP Global Forecast System forecasts at the same resolution. These files are independently downloaded for both integrations from the NCEP ftp server.

Once completed, both forecasts are uploaded to the BDFC main server, where products are generated from the default forecast. If this default forecast is not available, they are generated from the backup forecast.

Each 0.10º forecast yields a 909 MB NetCDF file and each 0.33º forecast a 38 MB NetCDF file. A GRIB2 file built from the forecast used to generate the products is archived. Its size is ~70 MB for the 0.10º forecast and ~7 MB for the 0.33º one.

The estimation of the current overall storage needs, including model data and products, as well as observations used for evaluation is of ~350 GB per year. Forecast images are daily sent by the BDFC main server to the WMO Global Telecommunications System through the AEMET operations service.

The BDFC main server is connected to the internet through a BSC Proxy/Firewall. The system architecture is drawn in Figure 2.

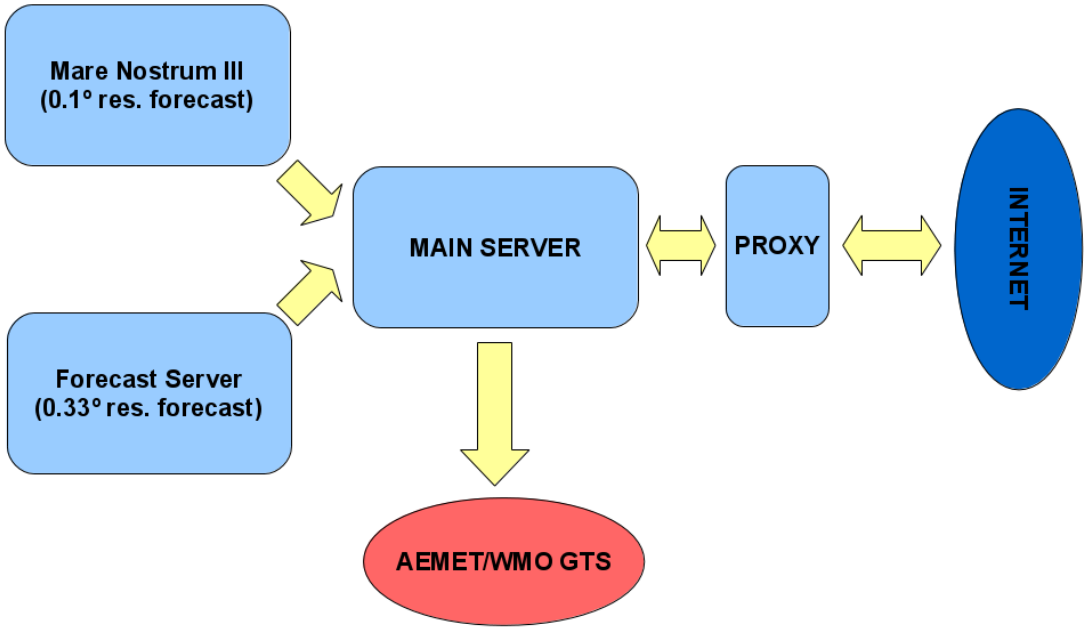

Figure 2 System architecture of the BDFC.

A system to notify eventual system issues has been implemented. It is based on:

- NAGIOS solution to monitor server and network services

- Email send to the technical staff in case any of the forecasts is unavailable
- Email send to the technical staff summarizing the daily operations

The three servers involved in the BDFC operations are the following:

I. BDFC main server

Software stack:

- Web server (Apache)
- Application server/Content Management System (Plone)
- Operating System: Linux - OpenSuSe Distribution

Hardware:

- CPU: total of 8 cores (Intel(R) Xeon(R) CPU E31230 @ 3.20GHz)
- cache size: 8MB
- Memory: 16GB
- Storage: 2TB
- Addressing: 64bit

Recovery:

- Daily incremental backup of the whole server (IBM Tivoli system) managed by the BSC operations team

II. Mare Nostrum III Supercomputer

Hardware:

- Peak Performance of 1,1 Petaflops
- 100.8 TB of main memory
- Main Nodes
  - 3,056 compute nodes
  - 2x Intel SandyBridge-EP E5-2670/1600 20M 8-core at 2.6 GH
  - 64 nodes with 8x16 GB DDR3-1600 DIMMS (8GB/core)
  - 64 nodes with 16x4 GB DDR3-1600 DIMMS (4GB/core)
  - 2880 nodes with 8x4 GB DDR3-1600 DIMMS (2GB/core)
- Xeon Phi Nodes
  - 42 heterogeneous compute nodes
  - 2x Intel SandyBridge-EP E5-2670/1600 20M 8-core at 2.6 GHz
  - 2x Xeon Phi 5110 P
  - 8x8GB DDR3-1600 DIMMS (4GB/core)
- 2 PB of disk storage

- Interconnection networks:
  - Infiniband FDR10
  - Gigabit Ethernet
- Operating System: Linux - SuSe Distribution
- MareNostrum has 52 racks and takes up a space of 120 m2
- Daily exclusive reservation of 260 cores to run the model

Recovery:

- Daily incremental backup of the whole server (IBM Tivoli system) managed by the BSC operations team

III. Forecast server

Software:

- Operating System: openSuSe Distribution
- Both nodes connected through Gigabit Ethernet using MPICHv2

Hardware:

- 2 nodes with 8 cores each (Intel(R) Xeon(R) CPU E5420 @ 2.50GHz)
- cache size: 6MB
- Memory: 16GB per node
- Storage: 2TB via NFS
- Addressing: 64bit

Recovery:

- Daily incremental backup of the whole server (IBM Tivoli system) managed by the BSC operations team

3. Risks and contingency plans

3.1. Known vulnerabilities

The most important vulnerabilities that have been detected are:

a. The BDFC main server has no redundancy.
b. The service depends on the resources (connectivity, power supply) of the BSC.
c. Both model runs depend on the same NCEP meteorological files used as initial and boundary conditions.
d. The configuration of the BDFC main server is not suitable to manage a large number of users and traffic.

3.2. Plans to reduce vulnerabilities

The following actions have been taken to reduce the vulnerabilities described in the previous Section.

a. Initial and boundary meteorological conditions

To reduce the dependency on the NCEP meteorological files, the BDFC plans to run the two model configurations with different initial and boundary conditions. The first tests have already be done to run the 0.10º forecast using meteorological files provided by the European Centre for Medium-Range Weather Forecast (ECMWF).

b. High reliability configuration

The risk of system failure will be minimized by duplicating the dedicated servers.

I. BDFC main server

- A duplicated 'slave' server, continuously synchronized with the 'master' will be configured at the BSC. User's requests will be redirected to the slave if the master is down.
- The possibility to duplicate the 'master-slave' server configuration at the AEMET Regional Center in Barcelona is under study. These servers would be activated when those in the BSC are down for electrical, network or hardware maintenance.

II. Forecast Server

- Similar actions will be taken as with the BDFC main Server.
- Although the 0.33º forecast will usually run with NCEP meteorological initial and boundary conditions, the possibility to automatically shift to ECMWF files in case of unavailability will be implemented.

c. High availability configuration

A high availability configuration is designed to manage a huge amount of users and traffic. Although, the BDFC is usually not in this situation, eventual access peaks, usually linked to tweet issuances, have been identified.

The configuration of the BDFC main server was modified in April after an outage linked to excessive concurrent accesses. Different configurations are currently under study to prevent further failures.