www.iblsoft.com

ibl

Web based NWP model management

Numeric Weather

Igor Andruska
Numeric Weather Team

Michal Weis
Managing Director

ECMWF Visualisation in Meteorology week 2015
15th Workshop on Meteorological Operational Systems
29th September 2015

Motivation - where do we come from

Numeric Weather

IBL develops variety of software for meteorological services:

- about 40% of ECMWF member/coop NMSs use IBL products operationally

Moving Weather

- Data collection & xfer
- Message switching
- forwarding to upstream systems

Visual Weather

- Forecaster's workplace
- Integrates a lot of data & tools to create fcst
- Visualisation
- OGC Web Services

Online Weather

- Web weather display
- Utilizing OGC Web Services
- Tools for forecasters
- Widgets for public web

Motivation - workflow & dataflow

Numeric Weather

Visual Weather - operational forecaster's workstation (but not limited to) - mainly integrates & displays observations, remote sensing data, and - models.

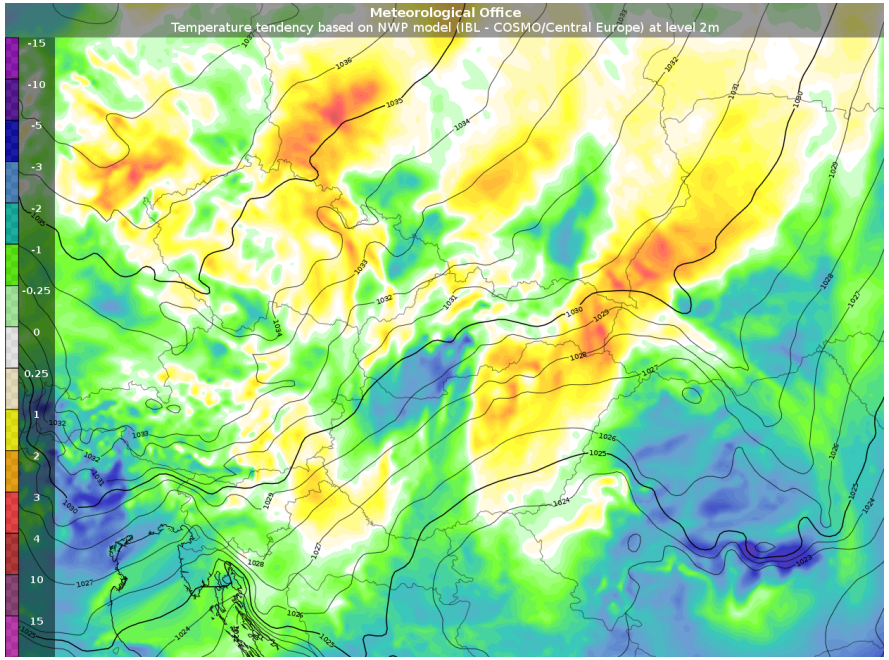

Used for: for operational forecasting, research, study/universities.

Consequently we deal with tasks of forecasters that depend on data (model) availability.

ibl

Numeric Weather

The “right” schedule

ibl

High-level design goals

Numeric Weather

- Build a NWP scheduler that can be operated and used without “guru-level” IT skills for:
  - ad hoc model runs by researchers, students ⇒ multi-user support
  - regular operational production
- NWP workflows created from predefined parameterizable functional blocks, just like Lego
  - Example of building blocks:
    - “Compute COSMO single domain forecast for initial time <T> and model domain <D>”
    - “Compute WAVEWATCH III forecast for initial time <T> and domain <D>, take driving wind fields from model <M>”

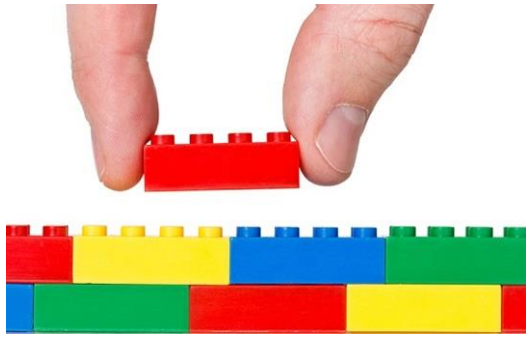

ibl

ecFlow strengths as we see it

Numeric Weather

- **Python API - the most important feature for IBL!**
  - Hard to live without API nowadays in general. End users often require unpredicted functionality
  - Allows us to implement our own extensions and facades on top of ecFlow core
- **Reliable - the most important for end-users (without much skills)**
  - Server is extremely stable (never crashed at IBL!)
  - Reliable handling of zombie jobs
  - Smooth recovery after power cuts
- **Support for multi-user environments**
  - It is straightforward to run a separate ecFlow server instance for each user without any undesired interference among users

ibl

ecFlow weaknesses #1

Numeric Weather

- **Missing possibility to parametrize suite in runtime without coding in Python**
  - Parameters: model initial time, forecast range, etc.
- **Built-in commands are bit low level (for non-daily users)**
  - To reliably stop a complex suite (kill all running jobs + prevent queued tasks from being executed) with complicated triggering (e.g., nodes triggered when other nodes abort) might require issuing a series of kills instead of just one. This is confusing for non-IT users
  - To run a node, users must distinguish between *begin*, *run, re-queue* and *restart* commands. The background is a little bit technical for non-IT users

ibl

ecFlow weaknesses #2

Numeric Weather

- **Lack of stable and user friendly UI for monitoring and control**
  - ecFlowview crashes a lot, complicated UI
  - Output of "*ecflow_client --get_state*" CLI command is hard to read
- **No intrinsic support for recursion** - (temporal) recursion is very common in NWP, e.g., for:
  - Continuous data assimilation
  - Ocean wave model restarted from the previous run

ibl

Numeric Weather

Building on top of ecFlow

ibl

Extended suite definitions

Numeric Weather

We extended ecFlow suite definition to allow runtime parametrization of tasks:

- Statements starting with #>
- Shell-like parameter expansion:
  - Environment variables !
  - Suite arguments %
  - Cross references to other parameters $
- Evaluated in runtime when task starts

```
#> domain = WW3_MediterraneanSea_125
#> wind-data = !{FORECASTS_ROOT}/GFS_World/%{run}/gfs-*
#> out-dir = !{FORECASTS_ROOT}/${domain}/%{run}
#> work-dir = ${out-dir}/workdir
#> prev-runs-dir = !{FORECASTS_ROOT}/${domain}

#> pbs-queue = batch
#> pbs-res = nodes=100%:ppn=12

#> ww3-rc = /opt/NumericWeather/rc/ww3_418.rc

suite ww3_single
    task recursion
        edit SCRIPT "load_and_run_suite"
        label def_file ""
        label run ""
        #> def-file = /home/nw/suites/ww3_single.def
        #> run = %{run-6}
        #> max-recursion-depth = 12 // 3 days
    endtask

    task workdir
        label work_dir "?"
        #> command = nw-ww3-create-workdir ${global:domain} ${global:wor
    endtask

    task ww3_grid
        trigger workdir==complete
        #> work-dir = ${global:work-dir}
        #> command = WW3RC=${global:ww3-rc} nw-ww3-grid
    endtask

    task initialize
        trigger ww3_grid==complete
        label restart_file "?"
        event cold_start
        #> work-dir = ${global:work-dir}
        #> command = WW3RC=${global:ww3-rc} nw-ww3-initialize -w ${globa
    endtask
```

ibl

Recursion

Numeric Weather

- Continuous model run: run initialized from forecast of the previous run
  - e.g. continuous data assimilation, warm start of ocean wave model, etc.
- We must make sure that the chain is always linked, no cold starts
  - e.g. by power outages, hardware failures, maintenance, etc.
- Need to model this relationship between successive model runs ⇒ recursion
- IBL solution:
  - special task in a suite (typically the first one) that
    - checks whether predecessor has been loaded in ecFlow. If no, loads it
    - waits until predecessor completes

```
suite cosmo_cda_2015092418
    task recursion
        edit SCRIPT "load_and_run_suite"
        label def_file ""
        label run ""
        #> def-file = /home/nw/suites/cosmo_cda.def
        #> run = %{run-6}
        #> max-recursion-depth = 12 // 3 days
    endtask # recursion
```

ibl

Macro commands

Numeric Weather

Using Python API, we combined elementary ecFlow commands into convenient macro commands with well defined behaviour that cover vast majority of use cases:

- **load-suite** - loads suite definition into server, parametrizes suite (e.g., model initial time) and (optionally) runs suite. All in single command
- **run-node** - runs suite/family/task regardless of its state (OK, there are few inevitable exceptions). If node is already running there is a switch to stop the node first and then run again
- **run-aborted-tasks** - runs only aborted tasks inside family or suite - common operation during recovery of the workflow
- **stop-node** - reliably stops suite/family/task in one shot regardless of state of node, its subnodes and triggering scheme
- **delete-suite** - deletes suite definition from server + removes all the files associated with the suite on the disk

ibl

Monitoring ecFlow

Command line interface

Numeric Weather

- With Python API we created a more user friendly version of *"ecflow_client --get_state"* command
- CLI switch to show/hide
  - triggers
  - labels
  - events
  - meters
  - variables
  - flags
- Use standard 16 terminal colors to highlight node state
- Useful for checking remote systems over slow internet lines

```
[nw@nw IBL]$ S ww3_single_2015092800 -lemt
suite ww3_single_2015092800 [active]
|   task recursion [complete]
|   |   label def_file = "/home/nw/suites/ww3_single.def"
|   |   label run = "2015092718"
|   task workdir [complete]
|   |   label work_dir = "/home/nw/forecasts/WW3_Mediterran
|   task ww3_grid [complete]
|   |   trigger workdir==complete
|   task initialize [complete]
|   |   trigger ww3_grid==complete
|   |   label restart_file = "/home/nw/forecasts/WW3_Medite
|   |   event cold_start : False
|   family wind [active]
|   |   trigger workdir==complete
|   |   task retrieve [active]
|   |   task ww3_prep [queued]
|   |   |   trigger retrieve==complete
|   family model [queued]
|   |   task ww3_shel [queued]
|   |   |   trigger ../initialize==complete and ../wind==co
|   |   |   label pbs_job_state = ""
|   |   |   label pbs_job_id = ""
|   |   |   label exit_status = ""
|   |   |   label eta = ""
|   |   |   event submitted : False
|   |   |   event running : False
|   |   |   meter progress = 0 <0;100>
|   |   task save_restart_files [queued]
|   |   |   trigger ww3_shel==complete
```

ibl

Monitoring ecFlow

Web-based interface

Numeric Weather

Suite	
cosmo_cda_2015092700	R
cosmo_cda_2015092618	R
cosmo_cda_2015092612	R
cosmo_cda_2015092606	R
cosmo_nested_with_cda_2015092618	R
cosmo_nested_with_cda_2015092612	R
wrf_relocable_2015092800	R
wrf_relocable_2015092712	R
wrf_slovakia_2015092800	Kill
wrf_slovakia_2015092718	Kill
▶ wrf_slovakia_2015092712	Kill
ww3_single_2015092800	Kill
ww3_single_2015092718	Kill
ww3_single_2015092712	R
ww3_single_2015092706	R

Node	Actions	State	Advanced
wrf_slovakia_2015092712	Kill	Show all details | Hide all details	
workdir #1	R Out	Click to show details	A
preprocess	R	Click to show details	
ungrib #1	R Out		A
metgrid #1	R Out	Click to show details	A
real #1	R Out	Click to show details	A
compute #1	Kill Out	**trigger** "preprocess==complete" **progress** 71/100 **submitted** **running** **pbs_job_state** "R: Job is running for 1h 1m 31s" **pbs_job_id** "39760.nw.iblsoft.com" **exit_status** "?" **eta** "0:24:53"	A C
postprocess		Click to show details	A C

```
2015-09-28 10:37:32,176|1049|INFO|***** JOB STARTED [ens_mem:0 try:1] *****
2015-09-28 10:37:32,456|1049|INFO|#!/bin/sh
2015-09-28 10:37:32,456|1049|INFO|
2015-09-28 10:37:32,456|1049|INFO|#PBS -N WRF_Slovakia_12-4km-wrf-metgrid-2015092806
2015-09-28 10:37:32,456|1049|INFO|#PBS -l nodes=4:ppn=12,walltime=02:00:00
2015-09-28 10:37:32,456|1049|INFO|#PBS -q batch
2015-09-28 10:37:32,456|1049|INFO|#PBS -m ae
2015-09-28 10:37:32,456|1049|INFO|#PBS -M nw-monitoring@iblsoft.com
2015-09-28 10:37:32,457|1049|INFO|#PBS -d /home/nw/forecasts/WRF_Slovakia_12-4km/20150
2015-09-28 10:37:32,457|1049|INFO|#PBS -o wrf-metgrid.out
2015-09-28 10:37:32,457|1049|INFO|#PBS -e wrf-metgrid.err
2015-09-28 10:37:32,457|1049|INFO|
2015-09-28 10:37:32,457|1049|INFO|nw-mpirun nw-wrf-metgrid
2015-09-28 10:37:32,457|1049|INFO|Submitting job using command: qsub "/home/nw/forecast
2015-09-28 10:37:32,547|1049|INFO|PBS job ID: 39766.nw.iblsoft.com
```

- Renders state of loaded suites
- Macro commands
- Zero-footprint installation

ibl

Monitoring ecFlow

Email alerts

Numeric Weather

1. Sent automatically when task aborts. Contains error message and Python traceback
2. Sent explicitly when certain conditions occur during processing, like missing crucial observation types for data assimilation

From nw@nw

Subject **[ecFlow][ERROR] /cosmo_nested_with_cda_2015072300/cda/obs/check_obs ABORTED**

To nw-monitoring@iblsoft.com

```
Job output tail:

2015-07-24 00:06:01,297|4085|INFO|***** JOB STARTED [ens_mem:0 try:1] *****
2015-07-24 00:06:01,329|4085|INFO|Observation directory: /home/nw/forecasts/COSMO_CentralEurope_0625_CDA/2015072218/workdir_cosmo/obs
2015-07-24 00:06:01,329|4085|INFO|Expected files: cdfin_synop,cdfin_synop_mob,cdfin_ship,cdfin_temp,cdfin_pilot_p
2015-07-24 00:06:01,329|4085|INFO|Mandatory files: cdfin_synop,cdfin_temp
2015-07-24 00:06:01,329|4085|INFO|Present files:
2015-07-24 00:06:01,329|4085|INFO|Missing files: cdfin_synop,cdfin_synop_mob,cdfin_ship,cdfin_temp,cdfin_pilot_p
2015-07-24 00:06:01,329|4085|INFO|Missing Mandatory files: cdfin_synop,cdfin_temp
2015-07-24 00:06:01,362|4085|ERROR|***** JOB ABORTED *****
2015-07-24 00:06:01,362|4085|ERROR|******************** Exception ***************************
2015-07-24 00:06:01,362|4085|ERROR|Traceback (most recent call last):
2015-07-24 00:06:01,362|4085|ERROR|  File "/home/nw/suites/cosmo_nested_with_cda_2015072300/cda/obs/check_obs.job1", line 114, in <module>
2015-07-24 00:06:01,362|4085|ERROR|    raise IBL.ECF.AbortError('Mandatory observation files missing')
2015-07-24 00:06:01,363|4085|ERROR|AbortError: Mandatory observation files missing
2015-07-24 00:06:01,363|4085|ERROR|******************** End of Exception ********************
```

ibl

Monitoring model outputs

Numeric Weather

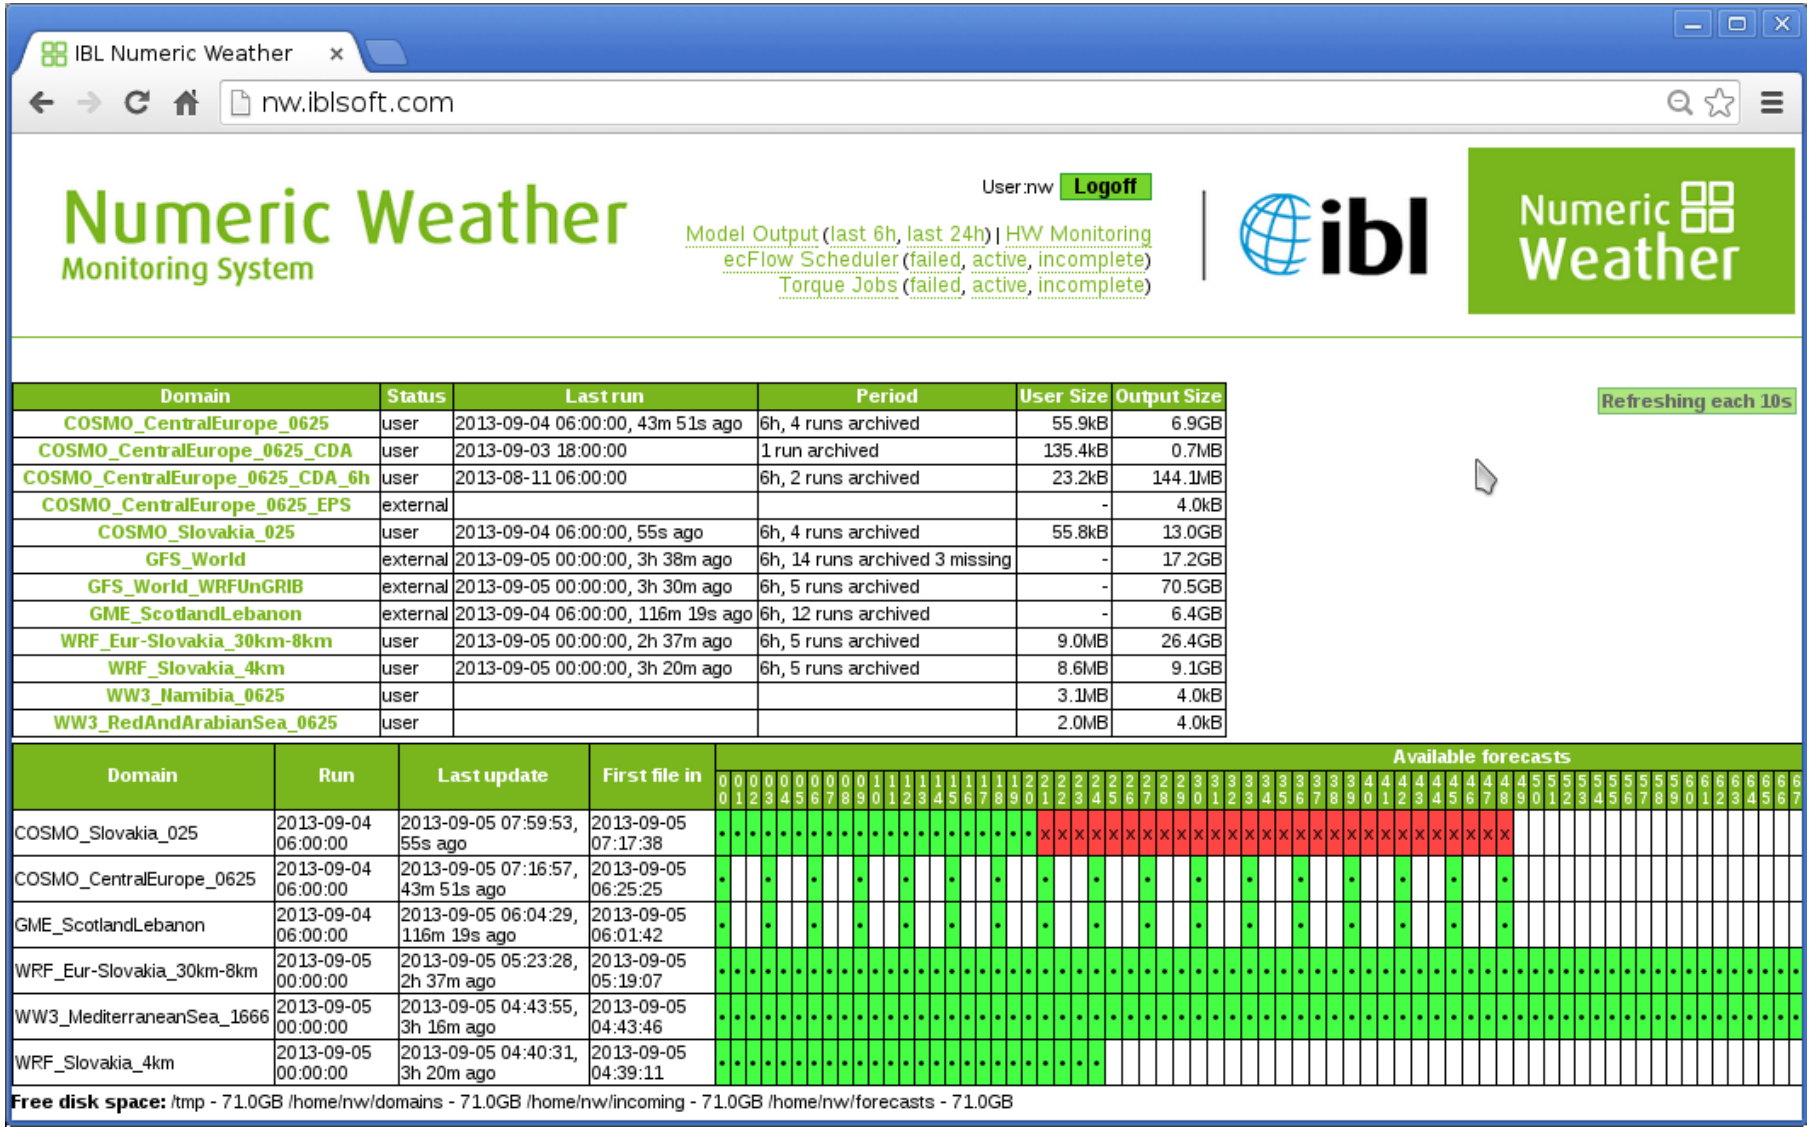

Domain	Status	Last run	Period	User Size	Output Size
COSMO_CentralEurope_0625	user	2013-09-04 06:00:00, 43m 51s ago	6h, 4 runs archived	55.9kB	6.9GB
COSMO_CentralEurope_0625_CDA	user	2013-09-03 18:00:00	1 run archived	135.4kB	0.7MB
COSMO_CentralEurope_0625_CDA_6h	user	2013-08-11 06:00:00	6h, 2 runs archived	23.2kB	144.1MB
COSMO_CentralEurope_0625_EPS	external			-	4.0kB
COSMO_Slovakia_025	user	2013-09-04 06:00:00, 55s ago	6h, 4 runs archived	55.8kB	13.0GB
GFS_World	external	2013-09-05 00:00:00, 3h 38m ago	6h, 14 runs archived 3 missing	-	17.2GB
GFS_World_WRFUnGRIB	external	2013-09-05 00:00:00, 3h 30m ago	6h, 5 runs archived	-	70.5GB
GME_ScotlandLebanon	external	2013-09-04 06:00:00, 116m 19s ago	6h, 12 runs archived	-	6.4GB
WRF_Eur-Slovakia_30km-8km	user	2013-09-05 00:00:00, 2h 37m ago	6h, 5 runs archived	9.0MB	26.4GB
WRF_Slovakia_4km	user	2013-09-05 00:00:00, 3h 20m ago	6h, 5 runs archived	8.6MB	9.1GB
WW3_Namibia_0625	user			3.1MB	4.0kB
WW3_RedAndArabianSea_0625	user			2.0MB	4.0kB

Domain	Run	Last update	First file in	Available forecasts
COSMO_Slovakia_025	2013-09-04 06:00:00	2013-09-05 07:59:53, 55s ago	2013-09-05 07:17:38	
COSMO_CentralEurope_0625	2013-09-04 06:00:00	2013-09-05 07:16:57, 43m 51s ago	2013-09-05 06:25:25	
GME_ScotlandLebanon	2013-09-04 06:00:00	2013-09-05 06:04:29, 116m 19s ago	2013-09-05 06:01:42	
WRF_Eur-Slovakia_30km-8km	2013-09-05 00:00:00	2013-09-05 05:23:28, 2h 37m ago	2013-09-05 05:19:07	
WW3_MediterraneanSea_1666	2013-09-05 00:00:00	2013-09-05 04:43:55, 3h 16m ago	2013-09-05 04:43:46	
WRF_Slovakia_4km	2013-09-05 00:00:00	2013-09-05 04:40:31, 3h 20m ago	2013-09-05 04:39:11	

Free disk space: /tmp - 71.0GB /home/nw/domains - 71.0GB /home/nw/incoming - 71.0GB /home/nw/forecasts - 71.0GB

Monitoring of hardware & system

Numeric Weather

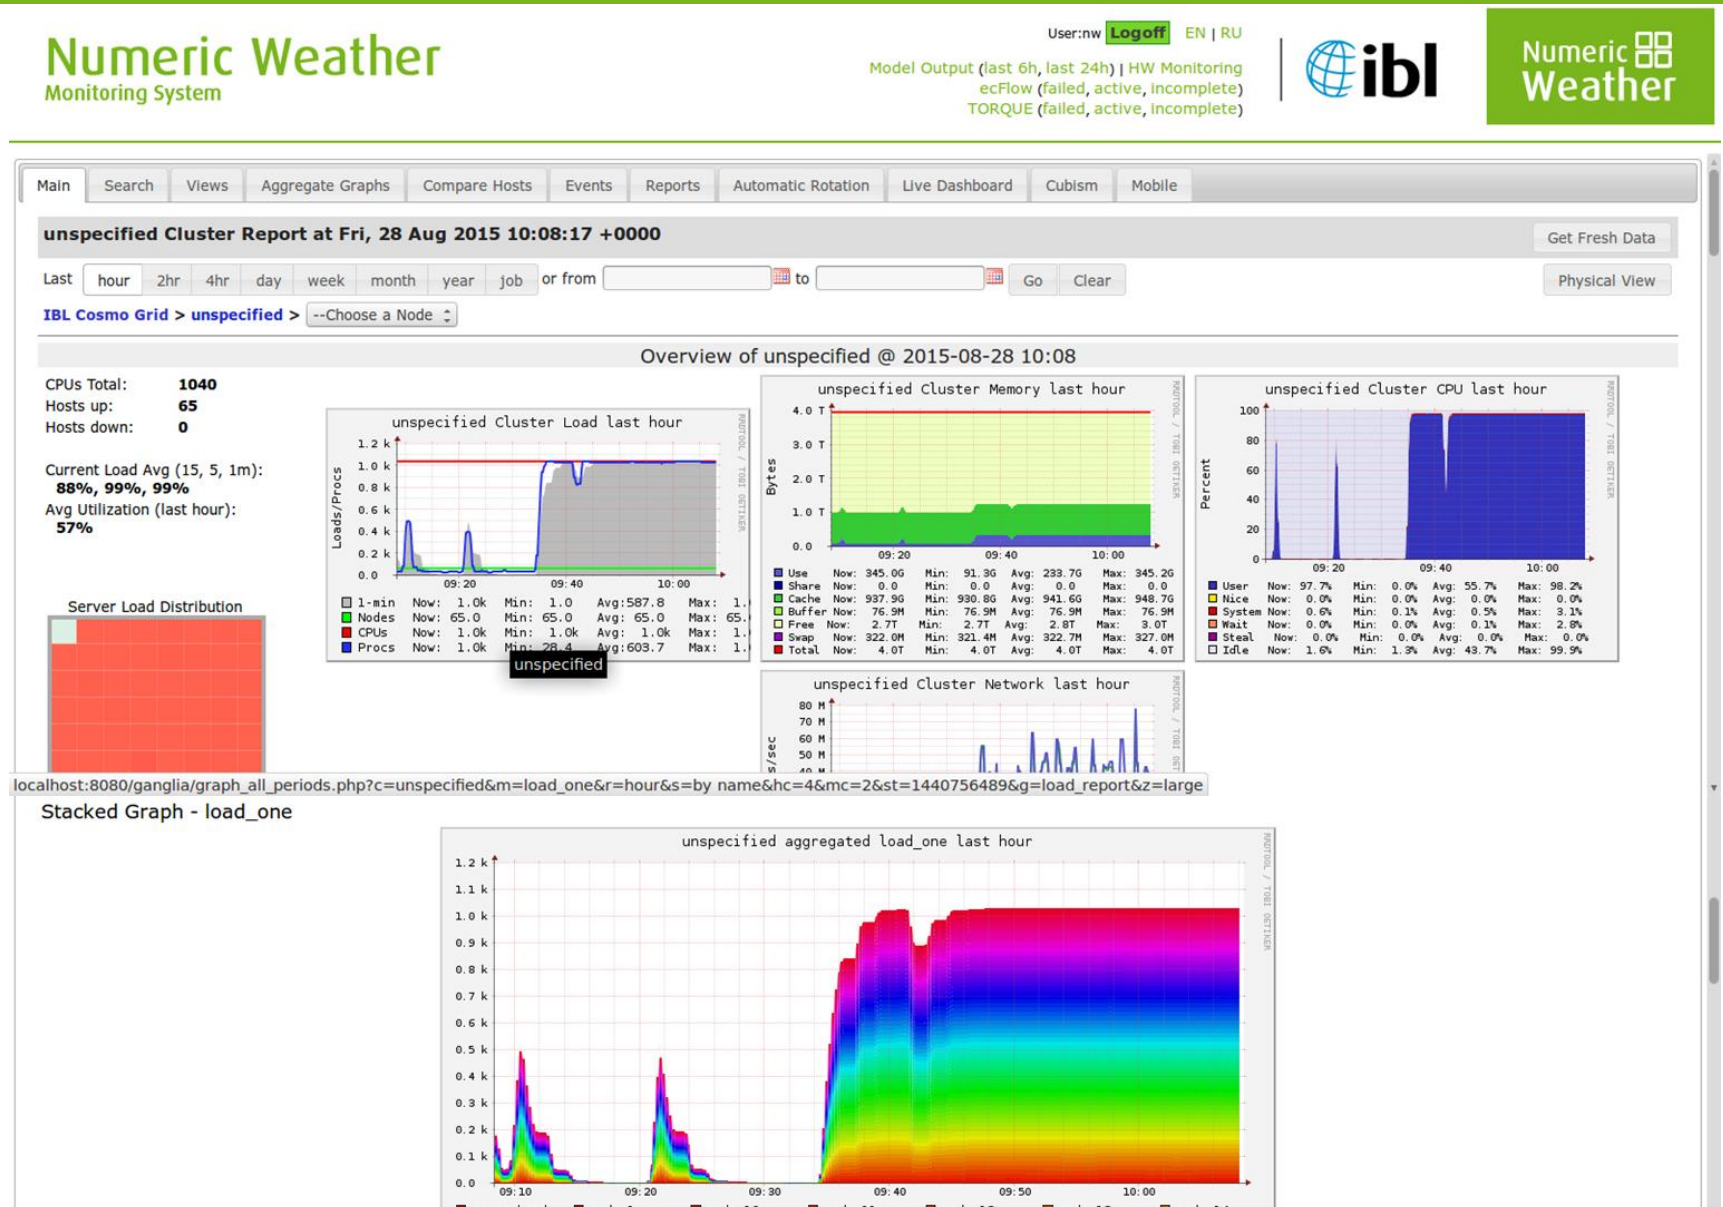

Numeric Weather

We will appreciate your comments and welcome further questions.

Michal.Weis@iblsoft.com • www.iblsoft.com

ibl