TECHNICAL MEMORANDUM

542

Report on the nineteenth meeting of Computing Representatives 22 - 24 May 2007

P. Prior (Compiler)

Operations Department

December 2007

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen terme

Series: Technical Memoranda

A full list of ECMWF Publications can be found on our web site under:
http://www.ecmwf.int/publications/

Contact: library@ecmwf.int

European Centre for Medium Range Weather Forecasts
Shinfield Park, Reading, RG2 9AX, England

Contents

Preface

The nineteenth meeting of Computing Representatives took place on 22–24 May 2007 at ECMWF. Twenty-two* Member States and Co-operating States, plus the CTBTO, EUMETSAT and JRC were represented. The list of attendees is given in Annex 1.

The Head of the Computer Division (Isabella Weger) opened the meeting and welcomed representatives. She gave a presentation on the current status of ECMWF's computer service and plans for its development. Each Computing Representative then gave a short presentation on their service and the use their staff make of ECMWF's computer facilities. There were also presentations from ECMWF staff members on various specific developments in the ECMWF systems. The full programme is given in Annex 2.

This report summarises each presentation. Part I contains ECMWF's contributions and general discussions. Part II contains Member States' and Co-operating States' contributions; all the reports were provided by the representatives themselves.

*Italy sent a report but did not attend the meeting.

Part I

ECMWF Staff contributions and general discussions

ECMWF Computing Service: Status and Plans – *Isabella Weger, Head of Computer Division*

Major activities over the past 12 months

High Performance Computing Facility

- The installation of HPCF Phase 4 has been completed.
- ITT 192 for the replacement of the HPCF has been prepared and issued.

Data Handling System

- The 2006 Phase of the DHS has been installed.
- A new Automated Tape Library has been installed in the DRS building.

ECgate

The Member States server ECgate is being replaced with a more powerful IBM p690 server to address performance issues.

RMDCN

The migration of the RMDCN to IP/VPN MPLS technology is underway.

Computer Division Organigram

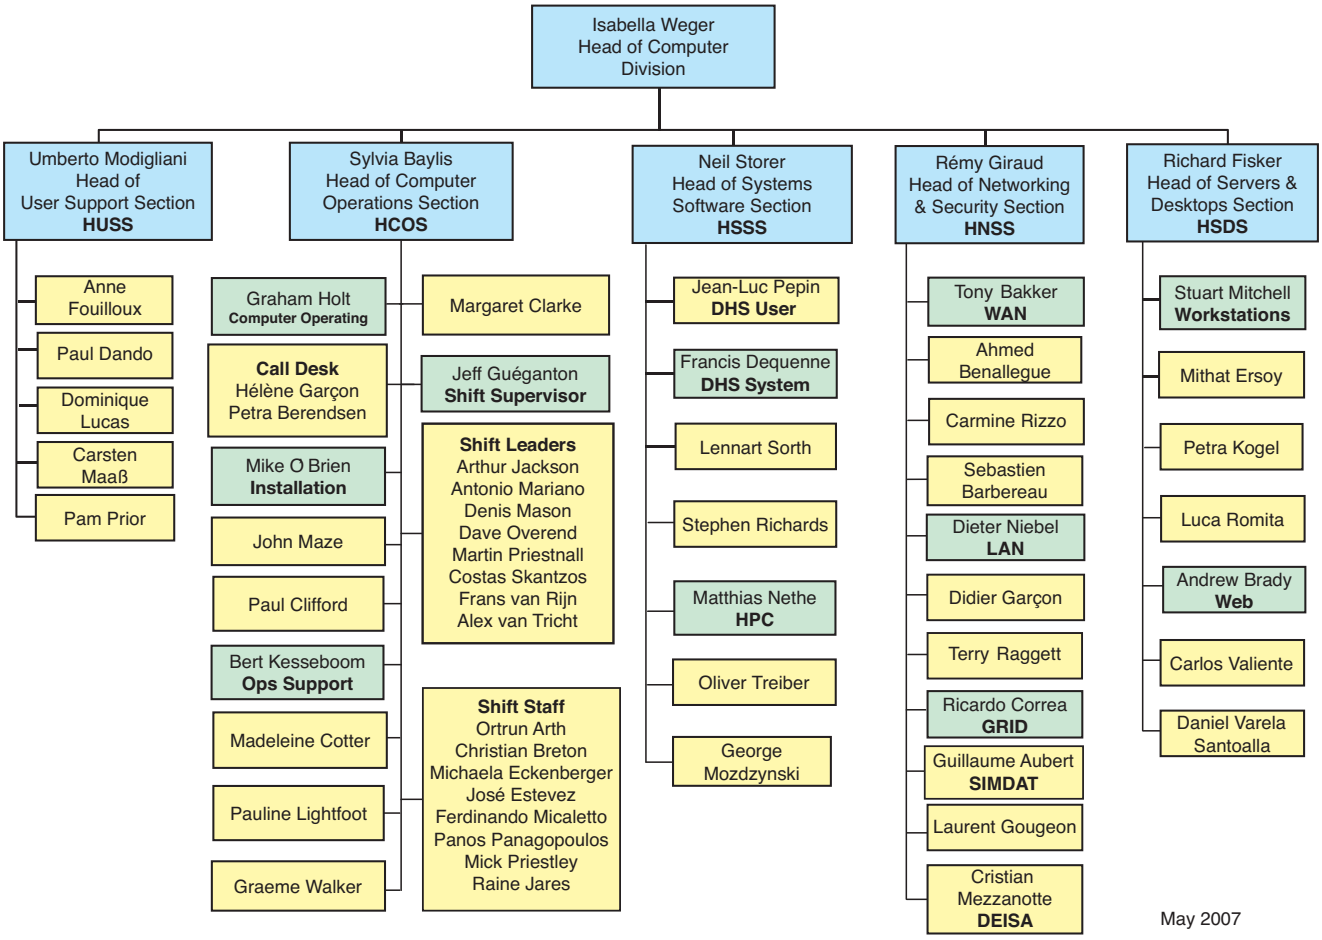

ECMWF Computer Environment

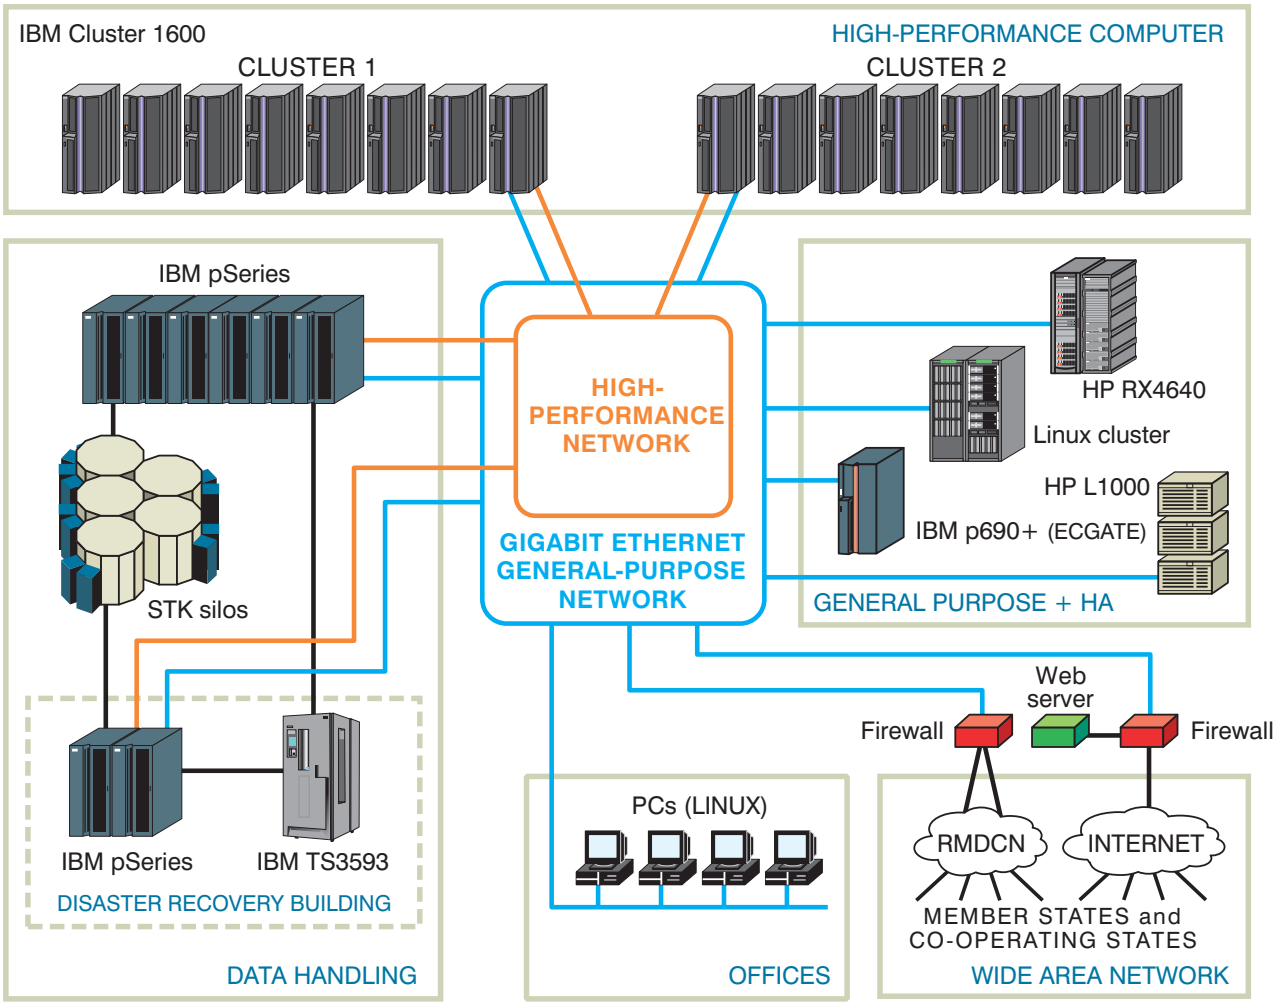

High Performance Computing Facility

The installation of Phase 4 of the current contract with IBM has been completed. The first cluster, "HPCE", was installed between March and mid July 2006. After the removal of HPCC, from the end of August until mid-November 2006, the second cluster, "HPCF", and MC-GPFS were installed. Early user access to the new hardware was possible in August 2006 and general access was available from September 2006. HPCD was taken out of service at the end of January 2007. In March 2007 the final nodes were installed. HPCE and HPCF now have 153 compute nodes each. Phase 4 passed its operational acceptance test in February 2007.

Current HPC system – IBM cluster 1600

There are two identical IBM AIX Clusters of p5-575 Servers.

Each cluster has 153 compute nodes (plus two "hot spares");

each node has:

- 16 Power5+ @ 1.9 GHz SMT processors;
- 32 GB memory (2 nodes with 128 GB);
- 10 nodes per cluster are dedicated to I/O and networking.

There is 50 TB of fibre channel disk storage per cluster.

The total sustained performance on ECMWF's codes is around 4 Tflops

ITT for the replacement of the HPCF

An Invitation To Tender (ITT) for the replacement of the HPCF has been prepared and issued. Following the Council session in December 2006, the Director established a Tender Evaluation Board. ITT documents were produced and released on 30th March 2007, with a closing date for tenders of 1 June 2007, after which the evaluation of the tenders will begin.

Timetable

30 Mar 2007	✓Issue of Invitation to Tender
1 June 2007	Closing date for receipt of tenders
Jul - Nov 2007	Evaluation of tenders and decision of the Director
Dec 2007	Signature of contract following approval by Council
Oct 2008	Start of parallel run
1 Apr 2009	Sole use of new HPCF system - Phase 1
2011	New HPCF system - Phase 2
31 Mar 2013	End of service period

HPCE weekly availability

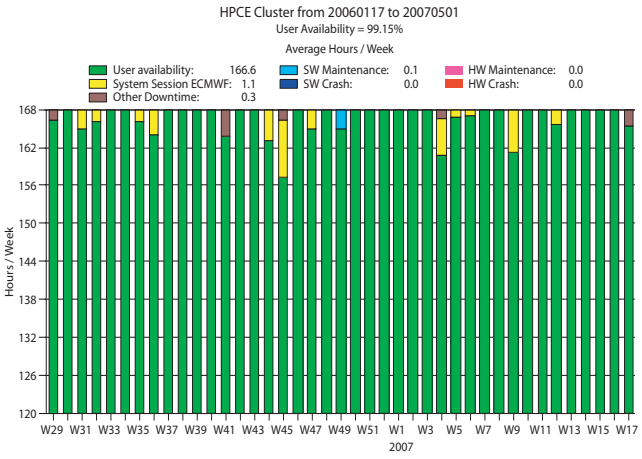

HPCF weekly availability

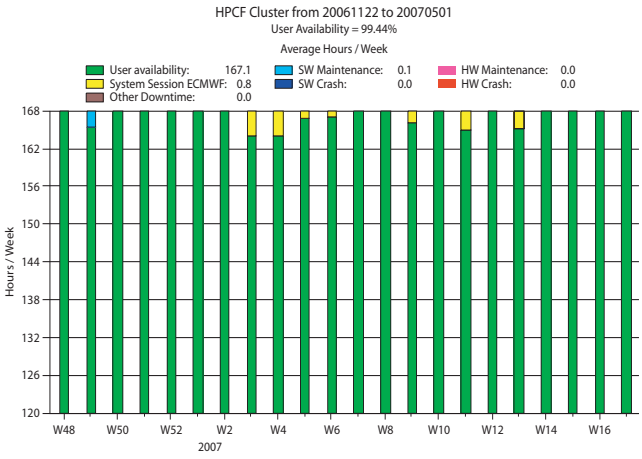

Usage of HPCF resources by Member States

Usage of HPCF resources by Member States has continued to increase. Due to changes in the overall availability of HPC resources in 2007 (additional nodes have been available since April and HPCD was still available in January), more HPCF resources will be distributed to Member States for 2007 than originally estimated. Early experience shows that the utilization of the new systems might be higher than that of the previous systems (above the assumed 90% utilization). The current estimate is that extra resources of between 8-13% of what is currently allocated could be added to the Member States' allocation.

The cost of the BC Optional Project may become higher than initially planned: the current estimate is for an extra 5% to be deducted from the participating countries.

Many Member State users have been assisted in migrating their jobs to the IBM Phase4 systems, in particular in use of the Simultaneous Multi-Threading (SMT) feature. All the Member State jobs submitted via the old SMS facility have been migrated to the new ECaccess-based mechanism. The old service has been terminated. The Framework for MS time-critical applications continues to be used for several applications.

Usage of NAG numerical libraries

ECMWF has started reviewing its in-house usage of the NAG numerical libraries which are available on both the servers and the supercomputers. The review has established that usage of the library is very limited, especially on the supercomputer. ECMWF is considering not installing the NAG library on the next HPCF.

Member States' representatives are asked for feedback:

- Do Member States' users rely heavily on the NAG library?
- Would you prefer us to inquire of all users directly?

Data Handling System

As of May 2007, 5.3PB of primary data are stored in the DHS; additionally, 2.7 TB of backup data in the DRS

HPSS 6.2 was installed in October last year. This has reduced overheads. New hardware and improved tuning has enabled MARS performance to be improved. Following an ITT for a replacement automated tape library in the DRS building, a new DRS robot and LTO-3 tape drives were selected and have been installed. The STK silos are reaching their end of life and will need to be replaced in the next few years.

DHS Archived Data (excluding backup copies)

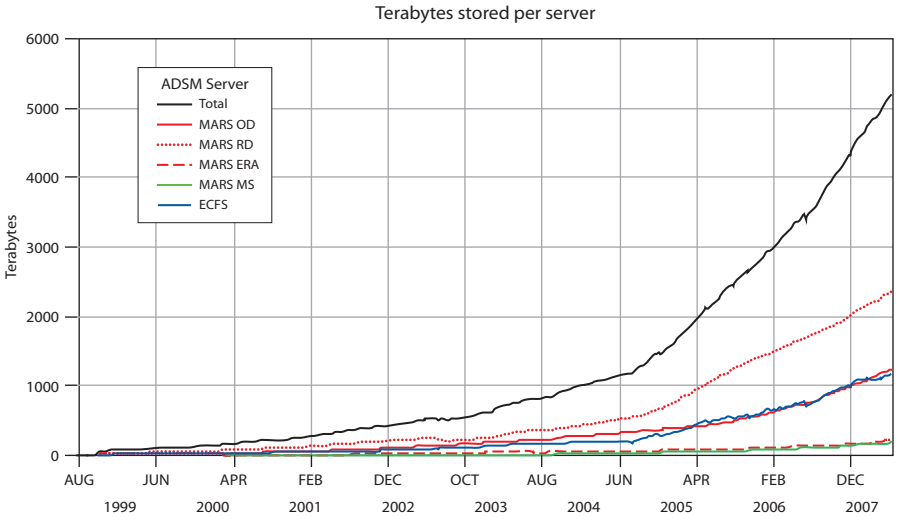

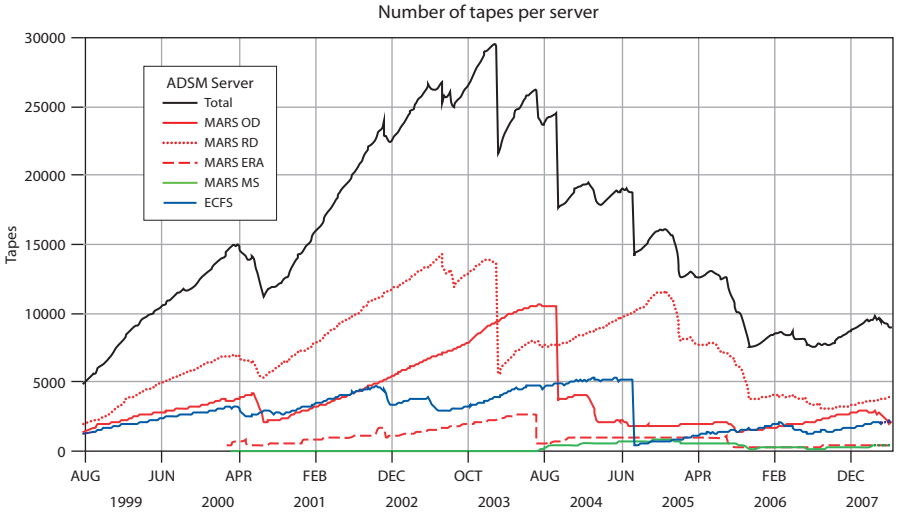

Servers and Desktops

The Member State server, ecgate, is being replaced by a more powerful p690 server. The HOME fileservers will be upgraded to new Itanium2 based systems, the same type of servers as are used for the HAPP data acquisition and dissemination system.

GPFS is being used on the internal Linux Cluster to serve SCRATCH filesystems and gain experience of it on Linux, as GPFS includes quotas, which are required for the management of the SCRATCH filesystems. (The installed Panasas filesystem does not support quotas in the way we require.) A Document Management system, Livelink from OpenText, has been installed; a pilot service is running for a small number of selected users. A new incident management system (Footprints) is being introduced into service.

Weekly availability statistics

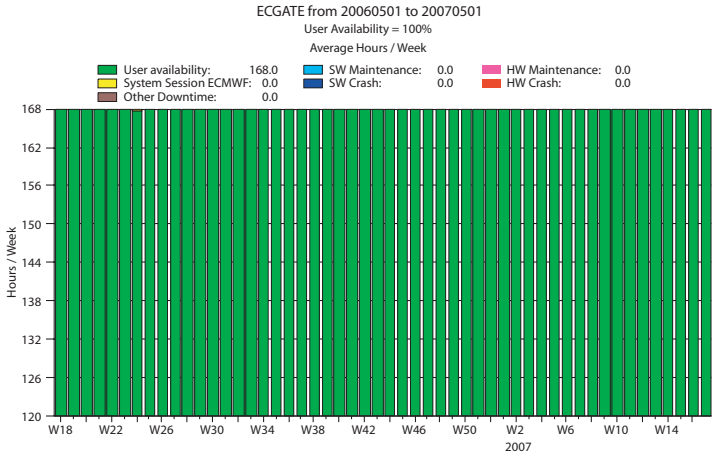

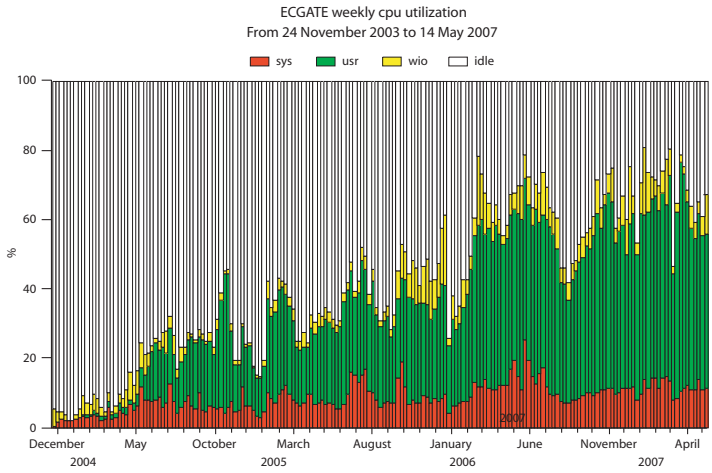

Web Services: Overview

The ECMWF web servers continue to provide a stable and reliable service. The average response time and availability of the web site remains excellent and use of the site continues to grow. The number of registered users accessing our site increased by 16% between 2005 and 2006 and preliminary figures for 2007 suggest that the 2006/2007 growth will be 13%. The average number of pages each of our registered users viewed per month increased by 4.6% between 2005 and 2006 and preliminary figures for 2006/2007 suggest a much stronger increase (16.2%).

Web Search facilities have been significantly improved. htdig search software has been replaced by Nutch/Lucene. A significant hardware upgrade was completed in January 2007; this has resulted in improved response, particularly from the popular forecast chart areas.

Web Services – Web Server Response

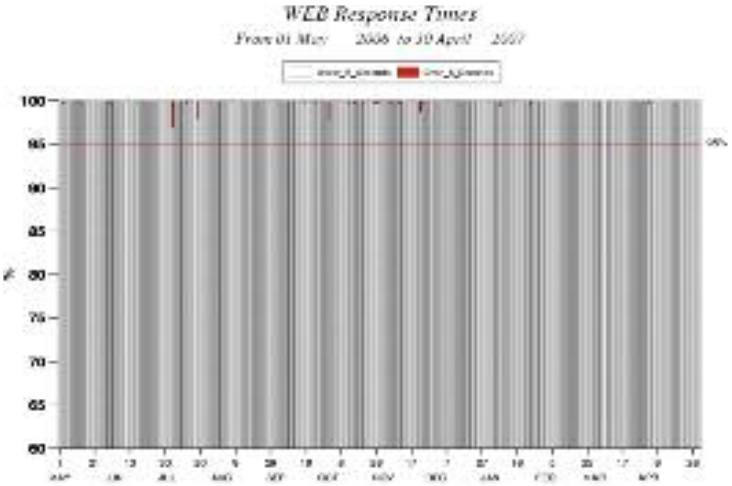

Web Services – page access

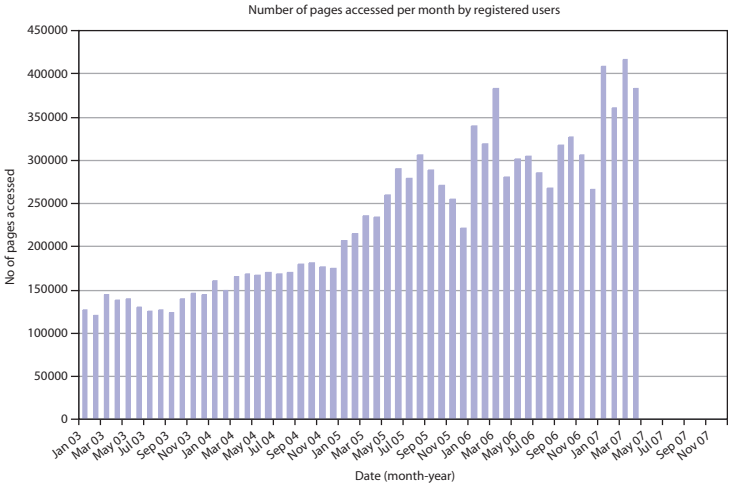

Web Services – Number of identified users

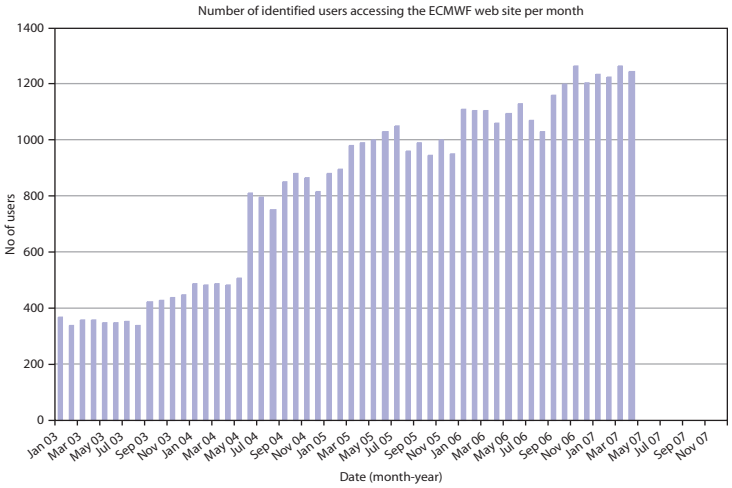

Web Service – Statistics

	2003	2004	2005	2006	2007*
Total page accesses by all users (millions of pages)	10.9	13.6	17.6	21.6	23.5
Change compared with previous year	+35.0%	+25.2%	+29.4%	+22.7%	+8.8%
Total page accesses by identified users (millions of pages)	1.56	2.02	3.05	3.69	4.70
Change compared with previous year	+68.7%	+26.5%	+51%	+21%	+27%
Average number of identified users accessing per month	-	685	974	1128	1240
Change in the number of identified users accessing the site	-	-	+42.2%	+15.8%	+13.3%
Average number of pages accessed per user per month	-	246	260	272	316
Change in the average number of pages accessed per user	-	-	+5.7%	+4.6%	+16.2%
Average time between page accesses (seconds)	2.89	2.31	1.79	1.46	1.34
Ratio of total users to identified users	6.8	6.8	5.7	5.9	5.0

**Based on first 4 months*

LAN

There was a major outage of the General Purpose Network in March 07. Late in the afternoon of 6th March, both core routers failed almost simultaneously. During the following night and the next day a temporary network was built to allow operational activities to resume. In parallel, the cause of the problem was analysed. A software bug, triggered by a configuration change the week before, was identified late on March 7th and normal operation resumed on the following day.

A new, Voice over IP-enabled telephone system, with a limited number of IP phones, has been installed. The migration will occur during the first weekend of June.

WAN

Frame Relay RMDCN

The current Frame Relay-based network, comprising 44 user sites, remains very stable and has a global availability consistently higher than its target of 99.5%. Its configuration has been frozen since the start of the process towards migration to IPVPN MPLS.

RMDCN Migration to IPVPN MPLS

The migration project started in late spring 2006. The Ready for Trial date was 16 April 2007, which was 3 months later than planned. User Site Acceptance began on 18 May 2007. The Reliability Acceptance Test is currently in progress. At present, Application Migration is scheduled for the first two weeks in June, starting on 4 June 2007, provided that the Reliability Acceptance Test has been passed.

Internet connection

A highly available, dual 250Mb/s connection to the ISP (UK Research Network SuperJANET) is being implemented. Improvements have been made to the Firewall and DMZ infrastructure.

Security

The development of the new Certificate Authority is complete. The migration of the various existing ECMWF Certificate Authorities is in progress. There is a new Strong Authentication system, based on ActivIdentity tokens. 900 tokens have already been deployed and all SecurID tokens will be replaced over the next 2 years. A review of the ECMWF Security Policy has started. It will be completely rewritten in compliance with ISO 17799 standard.

GRID activities – DEISA

ECMWF continues to participate in the EU funded DEISA (Distributed European Infrastructure for Supercomputing Applications) project. The former Phase 3 test cluster HPCU (6 nodes), is now deployed on a separate isolated network as a DEISA host and shares a common GPFS file system with partners in France, Italy and Germany.

The shared file system has proved to be more stable than anticipated, but operational issues remain to be addressed. Software to stage files from/to local storage to/from the global shared file system has been developed as part of the data federation activity. There have been no significant developments in the security area that would allow reconsideration of the decision to keep the DEISA host isolated from the operational clusters at ECMWF.

GRID activities – SIMDAT

ECMWF continues to co-ordinate the meteorological activity of the SIMDAT (Data grids for process and product development using numerical simulation and knowledge discovery) project. The V-GISC prototype was demonstrated at the WMO CBS in November 2006. The software is running continuously and catalogues are updated constantly. Data are available from 11 sites worldwide and more than 27,000 datasets are discoverable. The project is generating interest within the meteorological and other environmental communities. SIMDAT requirements are now used to define the requirements of the WIS and the SIMDAT infrastructure is seen as a major infrastructure for implementing the WIS.

Major ongoing/planned activities

- Complete the ITT process for replacement of the HPCF
- Continue to upgrade the DHS to ensure a balance between the power of the HPC system and the archiving and retrieval capabilities of the DHS
- Plan the replacement of the STK silos, which are reaching end of life and will need to be replaced in the next few years
- Continue to enhance ECFS to provide better performance and additional features for users
- Plan further improvements to the general purpose compute service for Member States
- Prepare an ITT to conclude a call-off contract (or call-off contracts) for PCs, laptops and small servers
- Complete the migration of the RMDCN from Frame Relay technology to MPLS (Multi Protocol Label Switching)
- Finalize the solution for Internet VPN backup for the RMDCN
- Plan the replacement of the High Performance Network (in time for the installation of the new HPCF)
- Rewrite the ECMWF security policy
- Replace one of the old chillers with a new, more efficient chiller
- Prepare an ITT for the enhancement of the electrical infrastructure in the computer building (main electrical power distribution)
- Complete the resolution of UPS issues
- Introduce an access security system in the computer building at the request of the insurer

T. Lorenzen asked why there had been a 5% increase in the computational costs of running the BC optional project. *Umberto Modigliani replied that the planned new cycle of the forecast suite, with a more refined radiation scheme and a higher resolution, had proved, during preliminary testing, to require more resources than had originally been estimated.*

R. Urrutia asked whether the Livelink document management system was open source or commercial and whether it was ITIL ready. *R. Fisker replied that it was a commercial product and that ITIL best practice does not cover document management. ECMWF's incident management system, Footprints, is ITIL ready but is not integrated with Livelink.*

R. Urrutia asked why Voice over IP (VoIP) had been introduced. *I. Weger replied that an exchange based on traditional telephony had been installed to replace obsolete equipment. VoIP capability was included in all new exchanges and ECMWF had taken the opportunity to test the facility with a limited number of IP telephones. It was not considered wise to have a telephone system dependent on its general purpose network. The prudence of this strategy was demonstrated by the incident in March when ECMWF's network was out of service for 36 hours.*

V. Gislason asked how long ECMWF had been using GPFS with Linux and what their experience had been. *R. Fisker replied that it had been under test since September 2006 and in use operationally for four weeks, although it was still being used via NFS for the time being. R. Fisker undertook to give V. Gislason more information, after further experience had been gained.*

Status of the Data Handling System – *Francis Dequenne*

The Data Handling System provides the medium and long-term archive of ECMWF. It is also used as a convenient transfer mechanism between ECMWF computers and between internal and external users. It is based on HPSS (High Performance Storage System) software. It can be accessed through two in-house applications: MARS and ECFS.

Volume of data stored

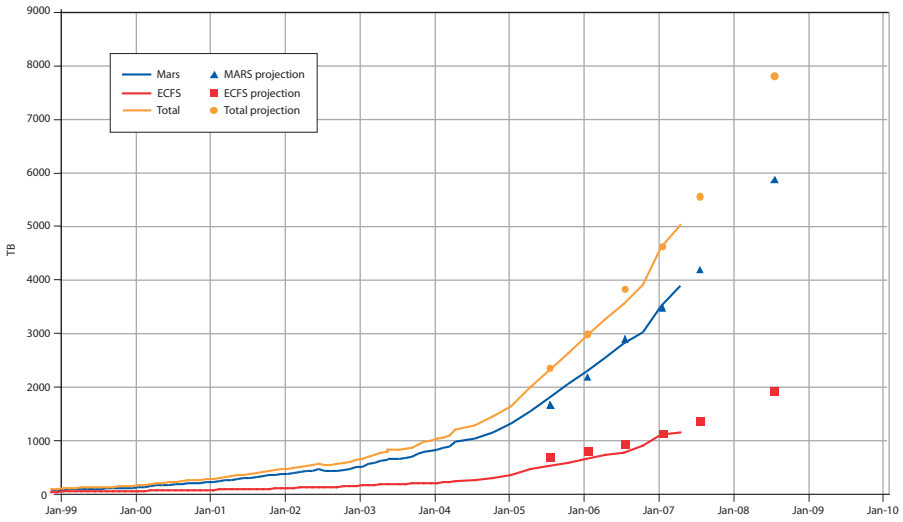

System usage

5.3 PB of data are stored, with an additional 2.7 PB of backups. There are 37 Million files. The archive grows daily by over 6TB.

A typical day's ECFS retrieval activity consists of about 30,000 files (80% from disk), that is, a volume of about 1 TB. Typical MARS daily retrieval activity comprises more than 9 million fields (70% from disk for MARS OD), that is, a volume of about 3.5 TB.

ECFS activity. Files retrieved/day

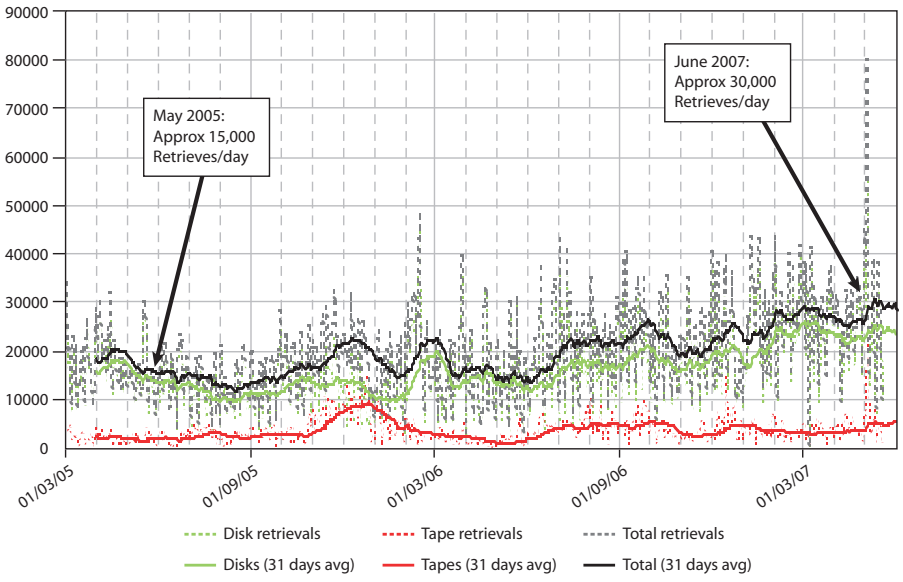

MARS requests

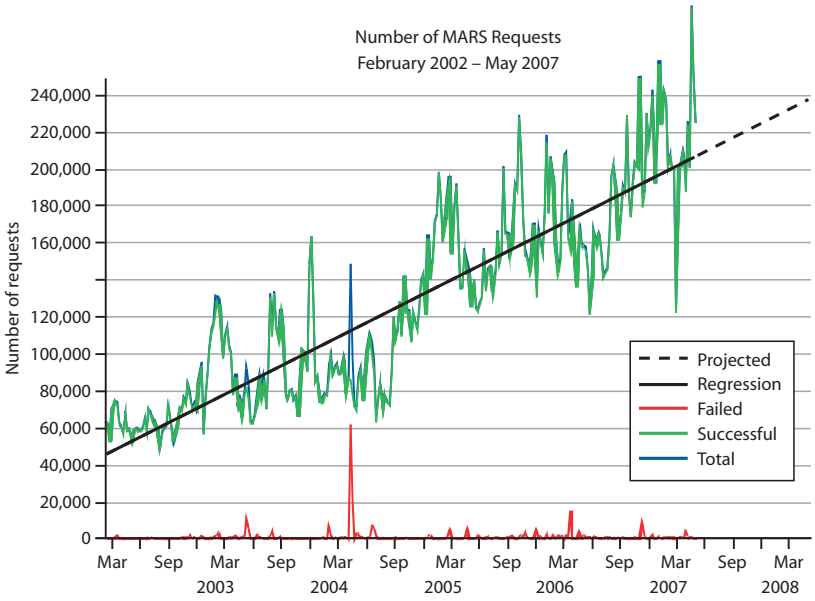

Data Handling System

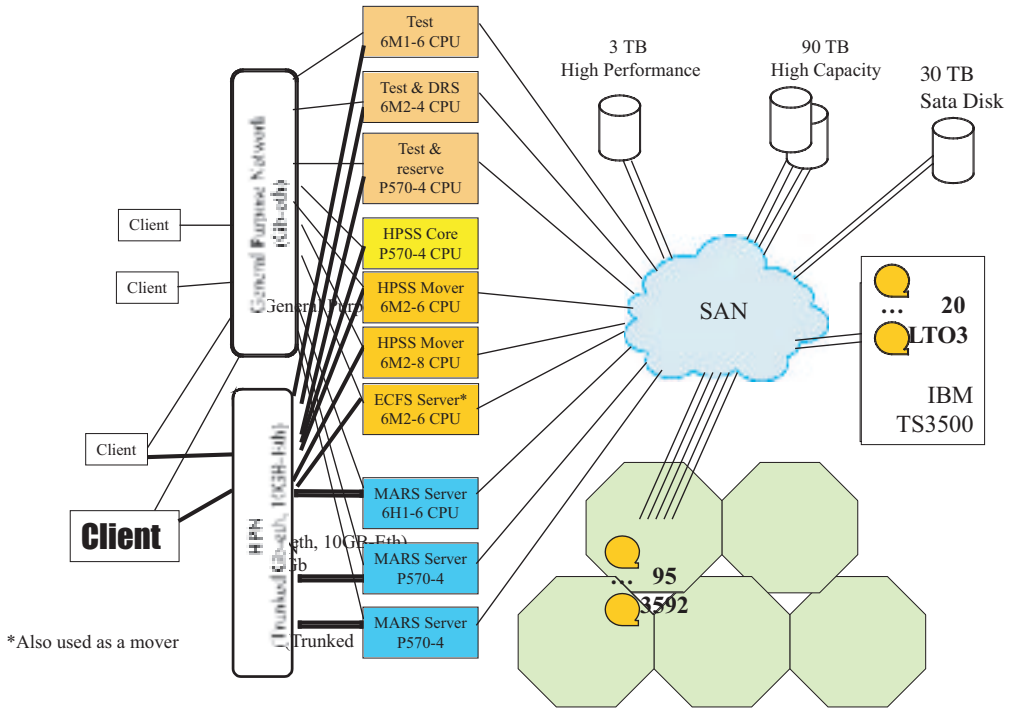

Recent developments

- HPSS 6.2. has been installed and DCE removed.
- Experiments with 3592 extended-length tapes have been carried out . These tapes allow 700GB to be stored on a 820m tape. They require 3592-E05 drives (also known as TS1120). They are now being used in production. Capacity will increase with later generations of the drives.
- The following new hardware has been installed:
  - New P570 servers have been phased in to replace older equipment.
  - New 4Gb-FC-able disk subsystems
  - Extensions to the 4Gb-FC SAN fabric
  - Most IBM 3592s were upgraded to E05.
- New DRS robot and tape drives were installed.

New DRS robot (Background)

The DRS robots are used to write backup copies of the most critical DHS data (currently 1.7PB of data); most of it is shelved after being written. The old ADIC AML/J is reaching its end of life.

An ITT was issued in 2006 to select a replacement. Valid tenders were received from IBM, SUN-STK, and ADIC. All bids proposed 20 IBM LTO-3 tape drives.

- IBM proposed a TS3500 1100 slot SCSI library
- SUN-STK proposed an SL8500 1448 slot library
- ADIC proposed a Scalar i2000 1000 slot library

All the bids satisfied ECMWF requirements. IBM provided the most cost-efficient, well-rounded solution and was selected.

New DRS robot (Status)

The new system was installed in early 2007. Some problems during the acceptance tests delayed operational deployment. These were recently resolved. It has been partitioned between TSM, our production and test HPSS. The TSM media and Test HPSS media have been fully migrated to the new robot. Production HPSS will be deployed shortly

Challenges encountered in the last 12 months

- Disk subsystem misbehaviours, which required disruptive microcode upgrades.
- SATA array performance issues: analysis and tuning of the SATA environment significantly increased performance.
- Deployment of the new supercomputers.
- Overall MARS retrieval performance: under heavy load, MARS was at times less responsive than hoped, so:
  - Better servers were provided to support the load.
  - A SAN reconfiguration was done.
  - More disk cache was allocated.

By and large, the quality of service has now improved. Occasional sub-standard performance is still observed, when reading data from tape.

- Staff shortage issues.

Future developments (short term)

- Deployment of "SAN-3P" facilities allowing HPSS movers to move data from disk to tape while limiting network data transfer.
- Test and possible deployment of newer generations of the IBM 3592 drives.
- Test of alternative drive technology.
- Version 7.1 of HPSS, allowing more efficient storage and movement of small files, especially on tape.

Future developments (long term).

- Support of the new High Performance Computer Facilities
  - The significant increase in capacity of the next HPCF will result in additional DHS data being stored in and retrieved from the DHS.
  - It is believed that the technical solution currently used will be adequate, when scaled.
  - Additional resources will be deployed in the DHS environment to cover this growth increase.
  - The STK silos will not be supported beyond December 2010
  - New robots, tape drives and possibly tapes will need to be acquired.
  - Phased replacement, possibly over 3 years.
  - ITT to be issued, possibly in early 2009.

R. Swennen asked why a SAN reconfiguration had been undertaken. *F. Dequenne replied that the way the disks were 'seen' by AIX was reconfigured, for instance larger blocks are used to access the SATA disks and 4 + 1 RAID arrays are used, instead of the previous 8 + 1 RAID arrays.*

P. Hitij asked what the SATA arrays were used for. *F. Dequenne replied that IBM SATA arrays were very reliable. Their performance was not so good as fibre channel disks (30-40 Mbps compared to 120 Mbps) but they were useful to store data which were unlikely to be needed on the HPC.*

Member State Server – ecgate, Status and Plans – *Richard Fisker*

ecgate status

It was originally planned to replace the current (old) ecgate by a Linux cluster. A major use of ecgate is for MARS data extraction, which uses interpolation heavily. Tests on the existing Linux Cluster showed that interpolation was slower than expected, based on performance running IFS, although using the latest processors would improve performance. The problem was discovered to be software based and has been fixed.

However, in the meantime, IBM agreed to leave one of the HPC Phase3 I/O nodes, together with the FAStT500 Fibre Channel Disk Subsystem. This provides approximately 3 times more processing power than the old ecgate and a significantly more capable I/O subsystem (I/O spread over 480 disks compared to 14). This was considered to be a better option for improving the service for Member State users by an effectively straightforward replacement - user work should move across to the new system with almost no changes required.

Mars GRIB Interpolation cpu times

System	Processor	CPU time seconds
"old" ecgate (AIX)	IBM Power4 1.3 GHz	53
Linux Cluster Node (old EMOSLIB)	AMD Opteron 2.2 GHz	104
Linux Cluster node (latest EMOSLIB)	AMD Opteron 2.2 GHz	60
"new" ecgate (AIX)	IBM Power4 1.9 GHz	30
"new" Linux node	Intel "Woodcrest" Xeon 5130 2.0 GHz	39

Currently, AIX systems use "Memory Mapped" I/O for coefficients.
Linux systems use "File I/O" ("Memory Mapped" was failing on Linux systems).
The above times are for post processing (GRIB extraction & interpolation) for a typical request for 492 fields over Europe to a 0.5 degree grid.

Old ecgate – IBM p690 with

- 16 1.3 GHz Power4 cpus
- 32 GB Memory
- 1 TB FastT700 disk subsystem with 2 controllers, using 14 x 73 GB disks

New ecgate – IBM p690 with

- 32 1.9 GHz Power4 cpus
- 132 GB Memory (124 GB usable for AIX & jobs)
- 8 TB FAStT500 disk subsystem with 24 controllers, 480 x 18 GB disks - 40 disks/controller. Fully redundant – no single point of failure
- Network connections to GPN (Single GigEth plus backup) and to DHS via HPN (trunked 4-way GigEth plus backup)

The new ecgate is running the latest version of AIX - 5.3, at the latest maintenance level available, plus the latest versions of compilers, libraries etc. The service on the new system took longer to provide than expected, since it was not at first possible to boot the system using all 32 cpus and all the Fibre Channel HBAs; this has now been corrected by a kernel patch (efix). The new system will be kept for at least 1 year.

A trial service will start on Wednesday 6 June, to allow users to test and migrate jobs; the MS SCRATCH filesystem will remain on the old system and be NFS-mounted on the new system; both interactive and batch access will be available. On 20 June we will make the new system the default "ecgate", and move the MS SCRATCH filesystem onto the new system.

We will keep the old system available for at least 2 more weeks, but this will only be for work which fails on the new system, to allow changes to be made.

It is still planned to install a Linux cluster to provide general purpose compute services for Member State users.

A New Strong Authentication Solution at ECMWF – *Didier Garçon*

The Evaluation Process - Main Criteria

- Security (all aspects)
- Ability to work in parallel with RSA SecurId
- Ease of integration with our Entity Management System (EMS), ECaccess, Web & Firewall
- Ease of development of the administration interfaces
- Ease of use for users and administrators
- Cost of the solution in the long term (10 years)

List of Vendors

- RSA
- Actividentity -Selected
- Vasco
- Verisign
- Secure Computing

Changes:

New tokens – Keychain V2 keypad tokens
Actividentity donated 2 special "big" tokens XL which have been assigned to 2 disabled users
New administration interfaces

Main differences from RSA tokens

Life expectancy

Tokens last 6 years (depending on usage).Tokens switch off their displays to save their batteries. They will not fail suddenly; close to end of life, they will show some signs - the display will get faint - but the token will continue to work for several weeks after this.

Tokens do not have any "expiration date". The first users to notice low batteries will have to request replacements. Then ECMWF will organise a whole new batch. Tokens running out of power are replaced free of charge by Actividentity.

Pin Management

The user's PIN is stored on the token. (With RSA the PIN is stored on the Server.) The correct PIN has to be entered on the token, in order to get a passcode. In other words, tokens can give only good passcodes (RSA tokens can give back a "wrong" passcode, if you enter a wrong PIN on the keypad.) This means that user errors are minimised: a wrong passcode can only be due to mistyping.

New tokens are programmed to have a default PIN when sent from ECMWF. After entering this default PIN, the user is obliged to change the PIN to a PIN of his choice. Tokens get locked (internally) after 10 incorrect attempts to enter a PIN. This is invisible to the Administrator, as it happens on the token. An administrative tool to unlock a token is available on the WEB Helpdesk: this involves a challenge response.

Account locked

- Accounts get locked ('disabled' in RSA terms) after 10 wrong passcodes. Administrators have full visibility of the number of wrong passcode attempts on the WEB Helpdesk interface.

The counter goes back to 0 when a good passcode is entered. Since the token always give good passcodes, this should not happen very often. The tool to unlock an account is also on the WEB Helpdesk

Synchronisation

Tokens have an internal clock (like RSA), and also an internal event counter that increments by one every time a password is obtained. Tokens can get out of synchronisation if this counter is out by more than 30 or the clock is out by more than 21 minutes. The tool to resynchronise tokens is on the WEB Helpdesk. Tokens have been programmed to wait for 3 seconds after a wrong passcode. This will discourage children from playing with the token!!

Synchronous/Asynchronous Modes

There are 2 modes in which a token can be used:

- Synchronous mode, also called Password mode
- Asynchronous mode, also called Challenge/Response mode or Dynamic mode

In Synchronous mode, the token displays a one time password of 8 digits straight after the PIN has been entered.

In Asynchronous mode, the token displays a prompt after the PIN has been entered to allow the user to input a challenge number displayed by the server. Pressing the red button will display the response to send back to the server. If no challenge is entered, pressing the red button will display a normal password.

User tokens have been programmed to support only the synchronous mode. Administrators' tokens have been programmed to support both modes, as the Authentication needed for the Administration WEB Help Desk can only be done by challenge/ response.

Display menu

When a passcode is displayed on the token, the menu button can be pressed to show more options

Documentation

Two leaflets have been created: one for users & one for administrators. Please make sure you include the correct one with every new token distributed.

Administration: Two interfaces

The Entity Management System Interface.

- Used for registering/deregistering users
- Will also be used to replace tokens (broken, lost, forgotten, stolen). This feature is being developed currently and will be available in June.

The Web Help-Desk Interface

- Used to list users and to fix any user problems (locks)
- Tool provided by Actividentity

Security Considerations and Awareness

- If a user calls you asking for his token to be unlocked, always make sure that you identify the person who is calling you as the real user.
- Be aware that anyone finding a token can lock the token, by entering 10 wrong PINs, and then ask an administrator by telephone to unlock the token (with a challenge/response) which will put the token in New PIN mode. This person (is he/she a genuine user?) can then gain access to the ECMWF account of the real user

Project status

- 700 ActivIdentity tokens were distributed in early April: a personal envelope, including token, documentation, and cover letter was sent to Administrators for each NMS user in their charge.

Tokens for users external to NMSs were sent to them directly, in order to minimise the workload for Administrators.

- A first version of the Token Administration documentation was sent to the main Administrators.

What remains to be done

- Options in EMS to replace broken, stolen, forgotten and lost tokens will be introduced -> being done currently, available early June.
- Administration documentation will be finalised, to include a PDF version available on the Web Interface -> also available in June.

E. Monreal observed that he had found the tool very straightforward to use but missed the "list of last access" feature on the long listing. *D. Garçon replied that they would investigate whether this facility could be provided.*

LAN: Overview & Resiliency Issues – *Dieter Niebel*

General-Purpose & High-Performance LAN

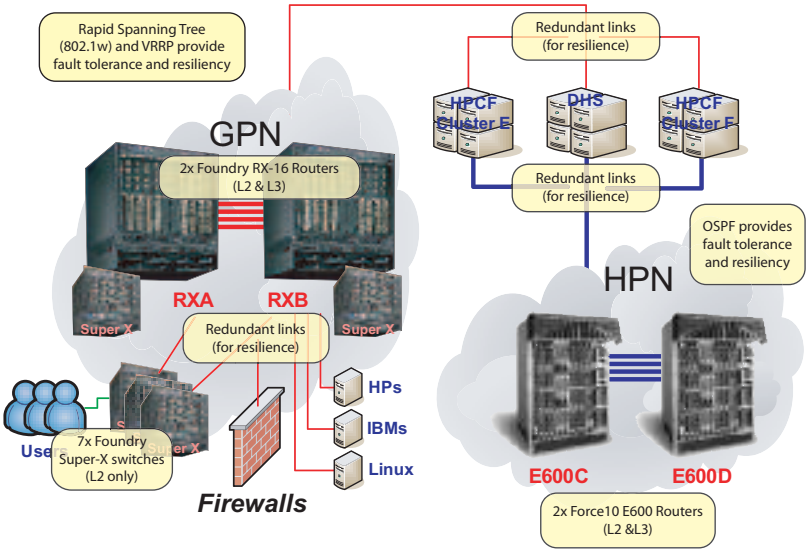

General-Purpose LAN on 6 March

Network outage began on 6 March late afternoon:

- RXA froze (this was noticed by the Nagios Event Monitor)
- After Power-cycling RXA, both routers were in a frozen state
- Further Power-cycling of either router made no difference

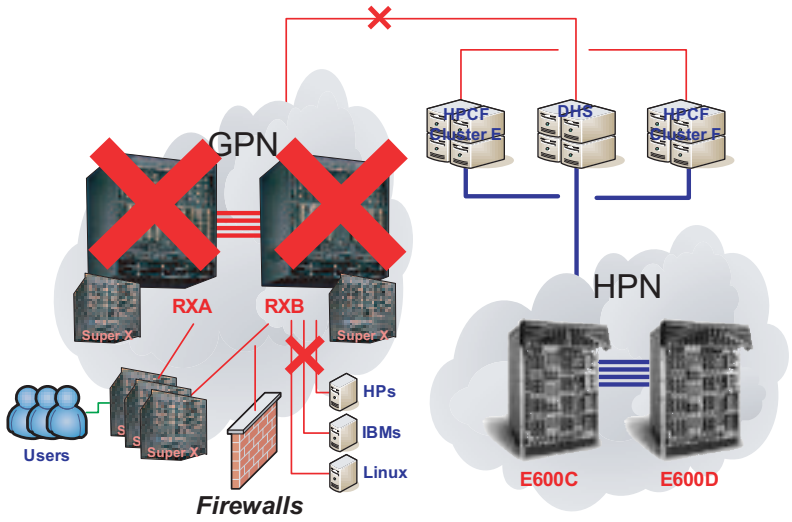

Workaround for Critical Operational Activities

- A temporary network was put together to allow critical operational activities to resume.

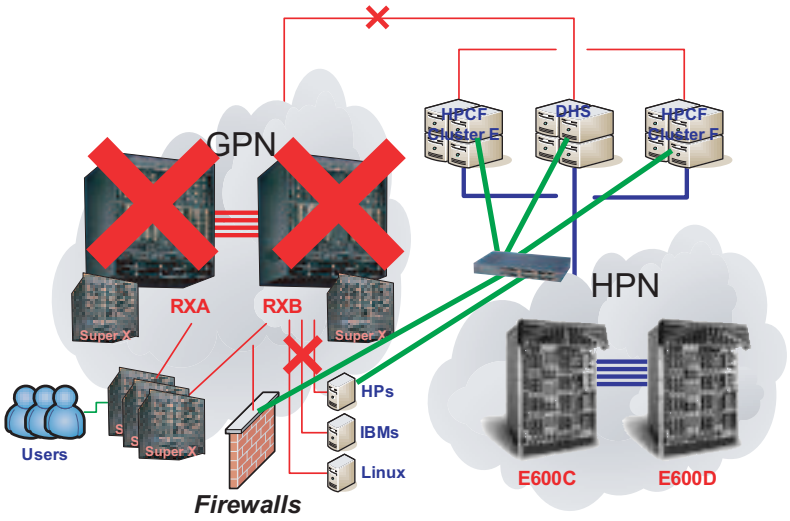

Workaround for Standard Operational Activities

- Another network allowed standard operational activities to resume
- A buffer overflow problem on the CORE router software and its trigger were identified as the cause of the outage after 24 hours
- Normal network service resumed at noon, 8 March

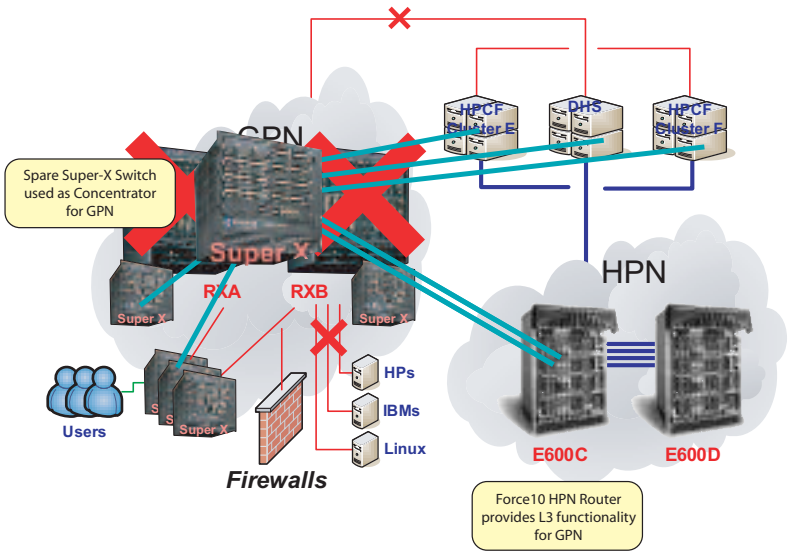

Nagios Event Monitor & CVS

- ECMWF uses a number of monitoring systems including Nagios, an Open Source system and service monitor for early detection of network problems. The Nagios monitoring daemon is set up to run regular checks for status and configuration changes on various network devices.
- The use of CVS (Concurrent Versions System) enabled us to access the current and previous versions of the configuration on the broken RX-16 Core Routers.

See http://nagios.org & http://www.nongnu.org/cvs

Remote Access during the event

- It became clear that Foundry Technical Support needed remote access to the Core Routers as quickly as possible.
- To provide this remote access, we
  - connected the consoles of the Core Routers to a Linux PC
  - connected the Linux PC to the Internet, and
  - temporarily configured the Internet router to allow remote access console of the RXB Core Router via SSH to the Linux PC.

Lessons

- The problem trigger turned out to be a configuration change done one week prior to the beginning of the Network Outage. A lot of time was lost in trying to find a hardware problem.
- A workaround for critical operational activities was put in place very quickly. This was because the problem started during working hours, when almost all analysts were readily available to help. It took a long time to provide Remote Access for Technical Support.

Lessons Learnt

- When considering possible causes, include configuration changes done since the last reboot of a system.
- Consider rebooting a system after a configuration change.
- To mitigate the impact of a similar scenario and to reinstate critical services quickly, in the event of a problem starting outside office hours, it is planned to:
  - Install a terminal server with Ethernet connections to the Internet Firewall and console connections to the Core Routers, so that remote access can be quickly and easily enabled.
  - Set up clearly identified management ports on all network devices, so that we can use a laptop as TFTP server to enable good configuration files to be downloaded.
  - Prepare an emergency network by setting up:
    - A sufficiently provisioned spare Super-X available next to the RXA Router, to act as concentrator
    - Cables ready to patch the spare Super-X to the HPN Router
    - Configuration changes to enable the HPN Router to provide L3 functionality to the GPN.

P. Halton asked: a) why the reconfiguration had been undertaken in the week prior to the network breakdown and b) whether it would have been possible to revert back to the pre-reconfiguration status. *a) D. Niebel replied that when a network node goes down on the HPC, the static route on the core routers should fail over from the failing node to the functioning node; this had not been happening and the reconfiguration had been carried out to remedy the problem. The 'route map' feature had to be applied on all the interfaces and the system could not cope with this.*

The suppliers, Foundry, are now developing a modification so that static routes can passively fail, if the node to which a particular route connects goes down.

b) D. Garçon replied the configuration is always saved before any change. They did revert to previous configurations but did not go back far enough; they should have gone to the configuration before the last reboot. The team did not consider the likelihood that a reconfiguration could cause problems which were not revealed until one week later, so they had followed incorrect assumptions.

R. Swennen asked why rebooting after a configuration change was recommended. *D. Niebel replied that this was purely a precautionary measure: if a reboot had been performed immediately after the configuration change, the problem would have been revealed immediately and it would have been obvious that a reversion to the previous configuration would remedy the situation.*

RMDCN Update – *Tony Bakker*

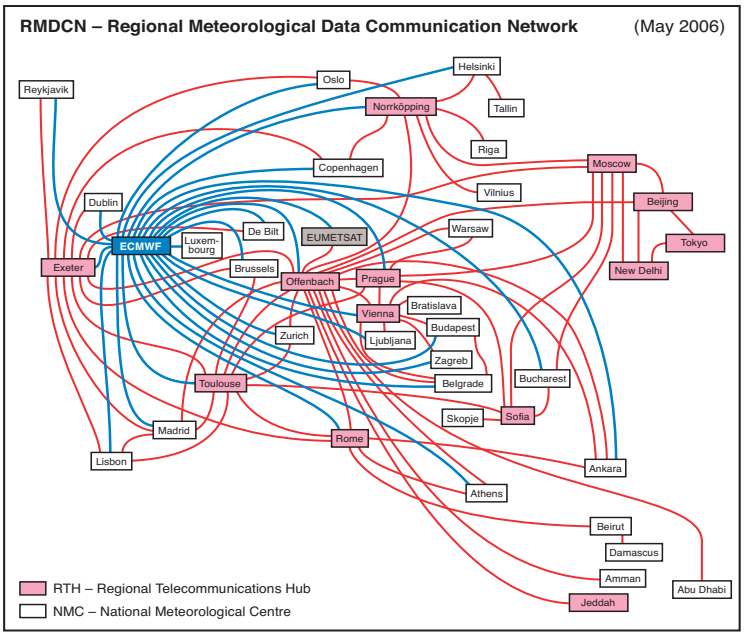

Current Frame Relay based RMDCN

- 44 sites (incl. ECMWF and EUMETSAT)
- Frozen configuration since Summer 2006

- Start of IPVPN MPLS Migration

- Global availability well above target (> 99.5%) for 2006 and up to April 2007

Migration to IPVPN MPLS

- Supplement no 4 of RMDCN Contract signed on 8 May 2006
- Final configuration and ordering completed in October 2006

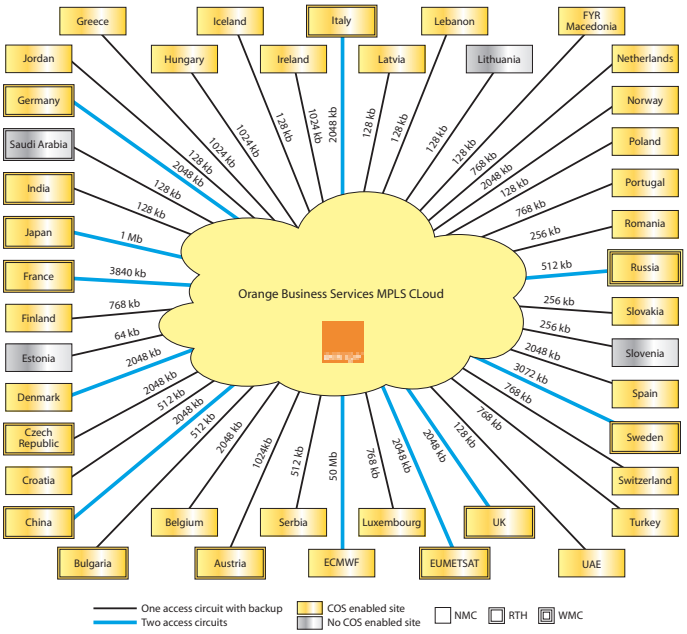

- Project schedule:
  - Initial Ready for Trial Date: 13 Dec 2006
  - First delay until 26 Jan 2007
  - Actual Ready for Trial Date: 16 April 2007
- User Site Acceptance Tests started
  - Was expected to last to two weeks, but lasted 5 weeks! – Network was not ready!

Summary of USA Tests problems

- Connectivity problems: Saudi Arabia, Poland
- Performance / Class of Service problems:
  - The IP Bandwidth values were wrongly configured (using the raw access line speed instead of the actual deliverable speed, eg. 2048 kbps instead of 1984 or 1920 kbps)
  - Some ACL misconfigurations (some User Sites test PCs in D1 class)
  - Routing issues (not all IP Networks were advertised)
- NAS Backup issues
  - Interfaces were in shutdown mode
  - HSRP was not tracking the correct interface
  - Incorrect switch type

Next Phases

- All User Sites were provisionally accepted on 18 May 2007. The Reliability Acceptance Test started on19 May 2007 and will run until 3 June 2007 - only 16 days!
- Application Migration Phase will start on Monday 4 June, provided that no serious problems arise during the RAT, otherwise it will be delayed until September 2007 (to avoid the summer holiday period).

Graphics Update – *Stephan Siemen*

Overview

- MAGICS
- Magics++
- Metview

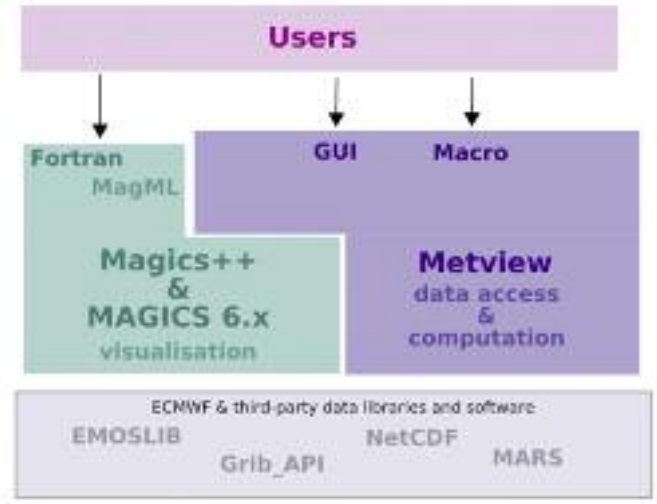

MAGICS

Maintenance of the current operational MAGICS has continued.

The changes for the switchover to the extended VarEPS were successful.

The latest internal test version is 6.12.

The current available export version is 6.11, released on 14 November 2006.

Platforms:

- Linux: SuSE 9.1
- SGI: IRIX 6.5
- Alpha: OSF1 V5.1
- IBM: AIX 5.1
- HP: HP-UX B.11.00
- Sun: SunOS 5.7

Magics++

The export version has been released to Member states.

The currently available export version is 2.0.1, released on 17 April 2007.

The latest internal test version is 2.1.0.

Platforms

- Linux: SuSE 9.1 / 10.x (32 and 64 bit)
- IBM: AIX 5

We would like to have feedback from Member States on which platforms (operating system / distributions + compilers) they use.

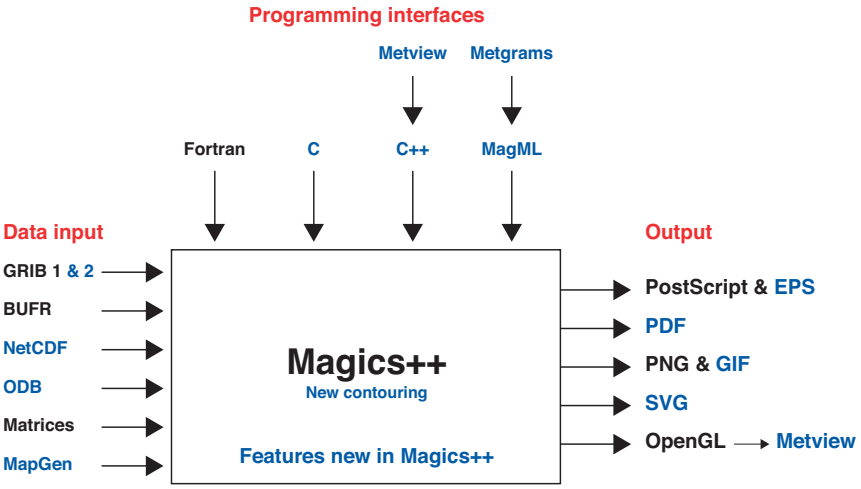

Magics++ – Grib 2

Used in monitoring TIGGE's Grib 2 data using MagML

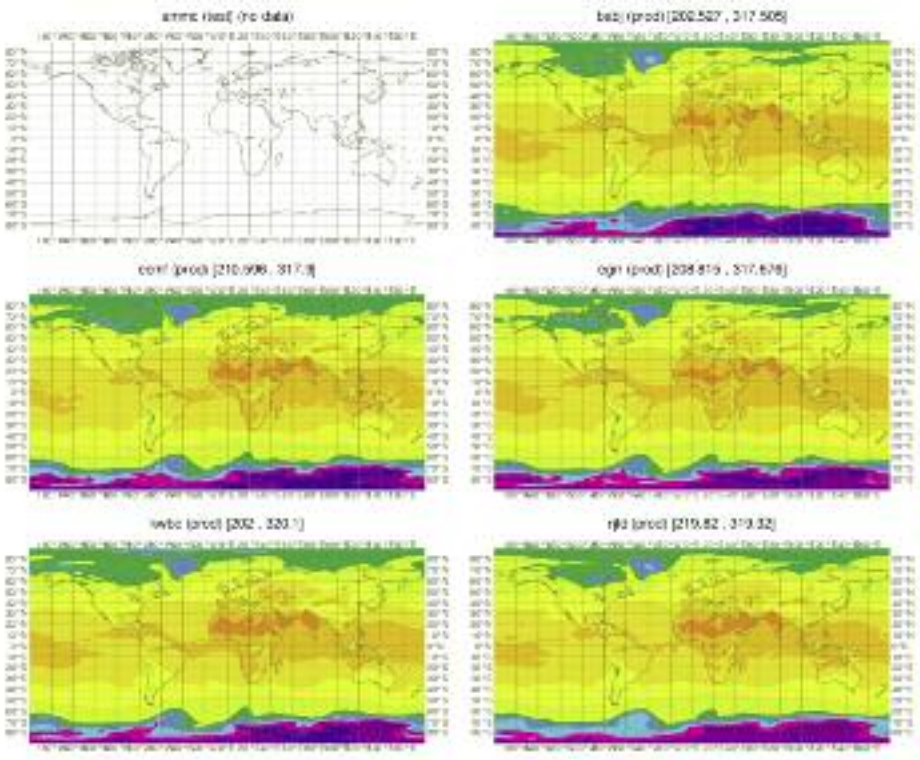

Magics++ Metgrams

Used to produce the new 15 day VarEPS metgram

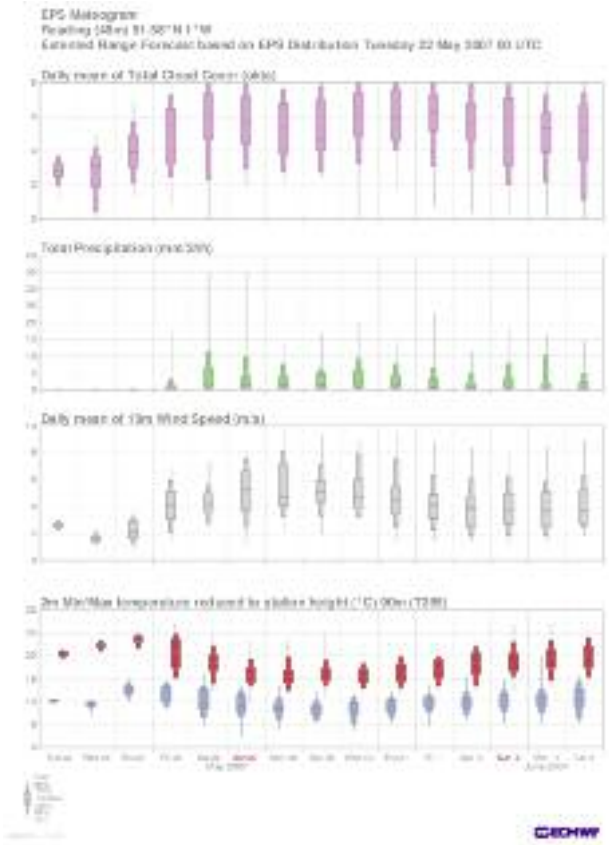

Magics++ – new features

- Boxplots available for users via an easy interface
- Maps with more than 360°in Longitude with wrap-around display of data
- Different resolutions for coastlines
- MapGen data files allow own borders and rivers to be added
- User logo

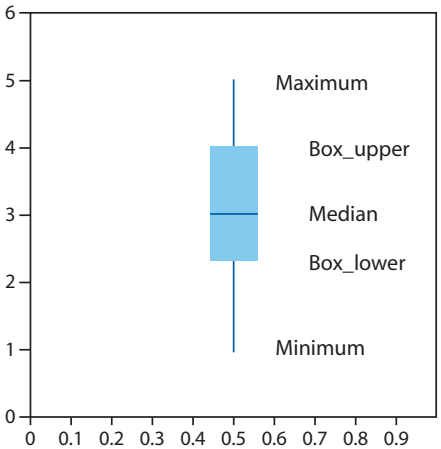

Magics++ – ecgate

- Available with the update of ecgate server
- First test version of Magics++ 2.1
- Delays in the release of an AIX version because of compiler problems
- Documentation on the webpage will be updated to show how to use Magics++ on ecgate.

Magics++ – plans

- Support for satellite
- New MagML syntax – taking into account all the feedback we received
- Basic observations
- Extensions to Metgrams, such as display of wind direction
- Improve support for NetCDF (3D data sets)

Metview – developments

- There is continuing routine effort to incorporate new facilities into Metview. These improvements include:
  - New requirements, such as enhancements for BUFR Edition 4
  - 'Percentile' application (EPS)
  - 'Height' and 'Expand' applications added to Hovmöller Family
  - Satellite Image re-projection application
  - EPS-metgram, 15 Days EPS metgram, and Classic Metgram now use Magics++
  - Enhanced Macro memory usage

Metview – Percentile

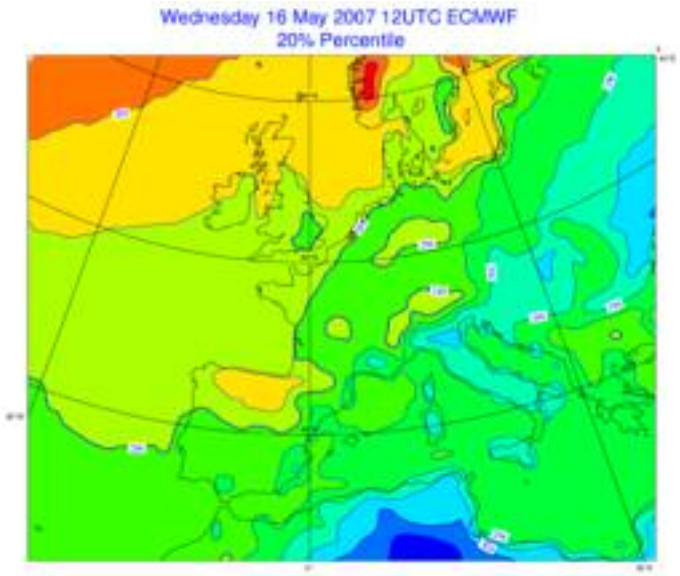

Metview – Hovmöller Diagram

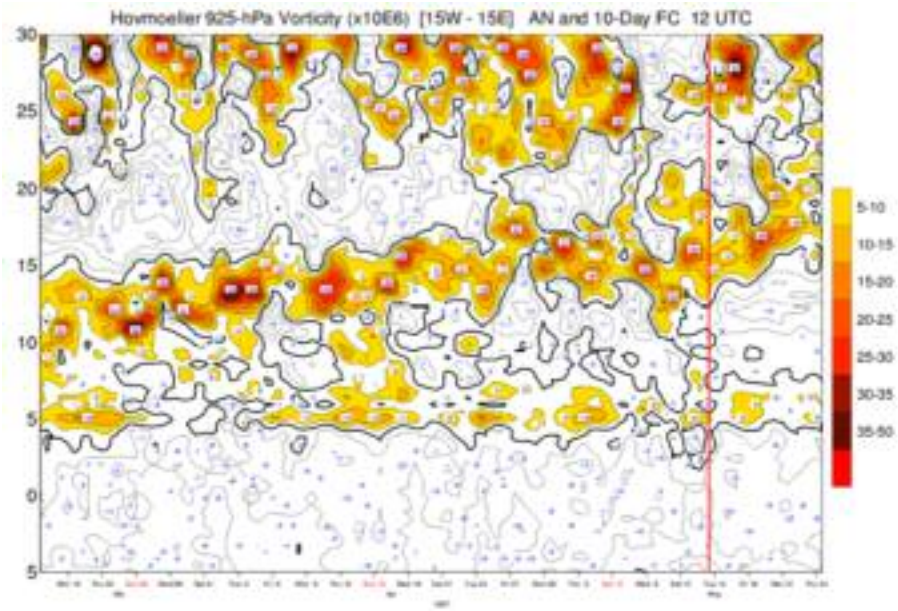

- New features were added to Metview to meet user requirements, including several Metview Macro functions:
  - remove_duplicates, var_a, stdev_a, covar_a, corr_a, search, tmpfile, frequencies, integer, unipressure, unithickness
- New compilers, such as GCC 4 and gfortran, were tested
- Routine efforts also include training courses, software maintenance and support for both internal and external Metview users.

Metview Macro – inline Fortran & C

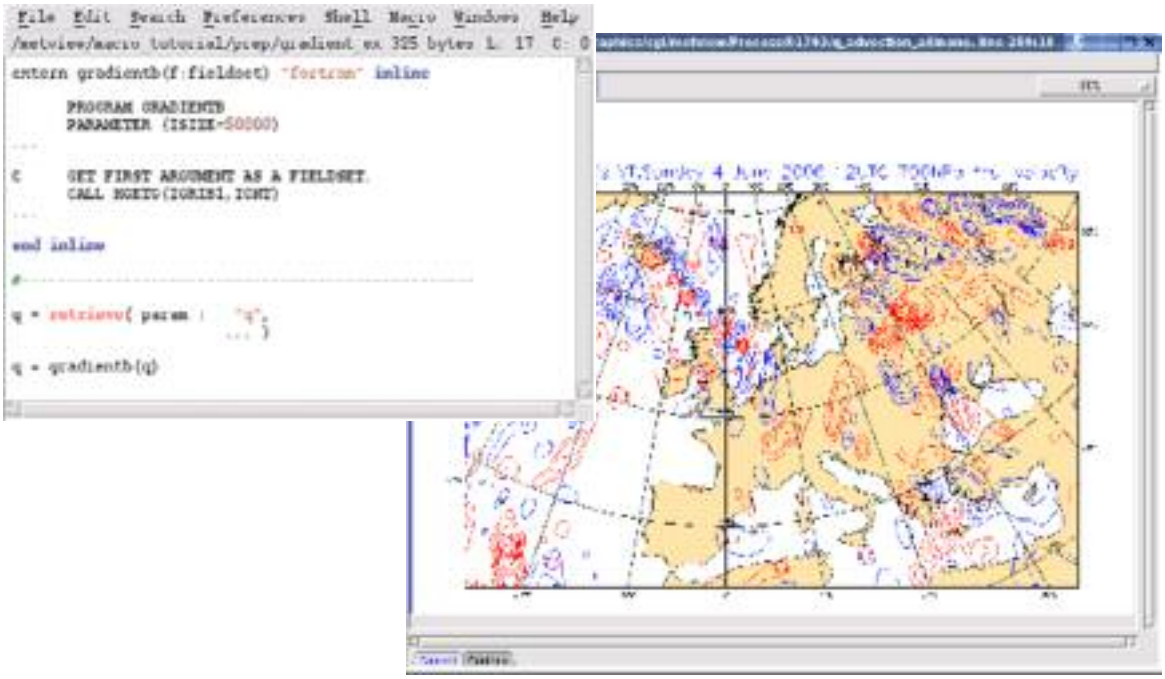

Metview – releases

- The latest internal Metview version is 3.9.3 which is based on Magics 6.12 and Emoslib cycle 000320
- Metview 3.9-export, released on 10 November 2006, supports the extended VarEPS
- Platforms
  - Linux: SuSE 9.1
  - IBM: AIX 5.1 / 5.3
  - SGI: IRIX 6.5
  - HP: HP-UX B.11.00
  - Sun: SunOS 5.7

Metview – plans

- Use of grib_api within Metview to handle Grib 2 data is being tested. This will be part of the next export version.
- A general purpose Macro library
- A new visualisation module to take advantage of all the benefits provided by Magics++
- The first prototype using Magics++ aims to execute current operational Metview Macros

Summary

- Metview
  - Metview 3.9 export version
  - Next export version 3.10 planned for 4Q2007
- MAGICS
  - MAGICS 6.11 export version
- Magics++
  - Magics++ 2.0.1 export version
  - Next Export version planned 3Q2007

For more information and to keep up-to-date with developments, please

visit our webpage or subscribe to our RSS at

http://www.ecmwf.int/publications/manuals/magics

http://www.ecmwf.int/publications/manuals/metview

P. Halton asked whether the recommendation was still to use SuSE Linux for Metview and Magics. *S. Siemen replied that Metview and Magics installation was currently being tested on various Linux distributions. He noted that an Appendix to the Metview Installation Guide contained notes on users' experience with various distributions of Linux. Moreover, he encouraged representatives to contact his section, if they wanted to test MAGICS and/or Metview on a new platform.*

R. Rudsar noted that they were having trouble installing the latest version of Metview on Fedora Core 5. *S. Siemen replied that their analysts should contact the Graphics section for advice.*

T. Lorenzen asked whether a graphical user interface to MagML existed. *S. Siemen replied that currently, standard text editors are used but it is planned to use something XML-based, as lots of programming languages have ways of supporting the writing of MagML. Some people use script languages to produce MagML output. It is not very user friendly but the idea is to write one script and then use it as a template for further work. Many users have reported that they find this method easier than writing and maintaining Fortran code.*

GRIB API update and new GRIB TOOLS – *Enrico Fucile*

GRIB API release version 1.0.0 (April 2007)

Encoding and decoding GRIB editions 1 & 2 is the main requirement and it is performed by the same functions for both editions (get and set). Due to the structural and semantic differences between the two coding systems, it was impossible to use a numeric table (as in gribex) to point to the information, as it is organised in a completely different way in the two editions. Instead, a Key/value based approach has been implemented: the information contained in a message is entirely described by a set of keys and their values, so that by using a few get and set functions it is possible to access and modify any value in the message, regardless of its edition or even its type.

There are two different types of keys:

- The coded keys are directly linked to octets of the GRIB message and their value is obtained by only decoding the octets.
- The computed keys are obtained by combining other keys (coded or computed) and when their value is set, all the related keys are set, in a cascade process. These keys provide a synthesis of the information contained in the GRIB message and are a safe way of setting complex attributes, such as the type of grid or the type of packing. Amongst these keys are: mars keys, angles in degrees, gridType, packingType, setDecimalPrecision ...
- A new set of command line tools is available, providing easy and reliable access to grib messages without writing code. With these tools it is possible to inspect the content of a grib file, modify the header of a message, copy some grib messages selectively from a file or compare two grib files.
- A GNU configure build system for installation is included in the distribution and has been tested on several platforms and several C and Fortran compilers. The build system is intended also to work on platforms which have not yet been tested. The only external library required is jasper for the jpeg2000 encoding and it is accessed automatically, if already installed in the system path. It is also possible to build without jasper, by disabling the jpeg2000 support.
- A set of tests is provided within the distribution to check that the library is working properly, after it has been built. It is recommended always to run the tests after each new build.
- A FORTRAN interface very similar to the native C interface is provided. Two identical sets of C and Fortran examples are included in the distribution, to show the similarity between the two interfaces.
- A manual is provided online: http://www.ecmwf.int/publications/manuals/grib_api/index.html and is included in the distribution. Some examples of how to use the tools are also included in the manual.
- The Library and tools are installed on all the platforms at ECMWF. The two environment variables, $GRIB_API_LIB and $GRIB_API_INCLUDE, are available to users for linking with the stable library version. The tools are available in all users' paths.
- grib_api can be downloaded from http://www.ecmwf.int/products/data/software/download/grib_api.html

GRIB Tools

- grib_dump
- grib_ls
- grib_get
- grib_copy
- grib_set
- grib_convert
- grib_filter
- grib_compare
- grib_get_data

grib_dump (WMO style)

grib_dump -HO file.grib

```
****** FILE: file.grib
======================   MESSAGE 1 ( length=10142 )           ======================
======================   SECTION_0 ( length=0, padding=0 )    ======================
1-4         identifier = GRIB
5-7         totalLength = 10142 ( 0x00 0x27 0x9E )
8           editionNumber = 1 ( 0x01 )
======================   SECTION_1 ( length=52, padding=0 )   ======================
1-3         section1Length = 52 ( 0x00 0x00 0x34 )
4           gribTablesVersionNo = 128 ( 0x80 )
5           identificationOfOriginatingGeneratingCentre = 98 ( 0x62 ) [European Center for
    Medium-Range Weather Forecasts (grib1/0.table) ]
6           generatingProcessIdentifier = 128 ( 0x80 )
7           gridDefinition = 255 ( 0xFF )
8           section1Flags = 128 [10000000]
9           indicatorOfParameter = 130 ( 0x82 ) [T Temperature K (grib1/2.98.128.table) ]
```

grib_dump (Debug)

grib_dump -D file.grib

```
=======> section GRIB (10142,10142,0)
   0-0 constant grib1divider = 1000
   0-0 constant ieeeFloats = 0
   0-0 transient dummy = 1
   =======> section section_0 (0,0,0)
      ----> label empty
   <====== section section_0
   0-4 ascii identifier = GRIB
   4-7 g1_message_length totalLength = 10142
   7-8 unsigned editionNumber = 1 [ls.edition]
   =======> section section_1 (52,52,0)
      8-8 constant ECMWF = 98
      8-8 position offsetSection1 = 8
```

grib_dump (C code)

grib_dump -C file.grib

```
#include <grib_api.h>
/* This code was generated automatically */
int main(int argc,const char** argv)
{
    grib_handle *h      = NULL;
    size_t size         = 0;
    double* v           = NULL;
    FILE* f             = NULL;
    const char* p       = NULL;
    const void* buffer = NULL;

    if(argc != 2) {
       fprintf(stderr,"usage: %s out\n",argv[0]);
       exit(1);
    }

    h = grib_handle_new_from_template(NULL,"GRIB1");
    if(!h) {
        fprintf(stderr,"Cannot create grib handle\n");
        exit(1);
    }

    /* empty */

    GRIB_CHECK(grib_set_long(h,"editionNumber",1),0);
    GRIB_CHECK(grib_set_long(h,"gribTablesVersionNo",128),0);

    /* 98 - European Center for Medium-Range Weather Forecasts (grib1/0.table)  */
    GRIB_CHECK(grib_set_long(h,"identificationOfOriginatingGeneratingCentre",98),0);

    GRIB_CHECK(grib_set_long(h,"generatingProcessIdentifier",128),0);
    GRIB_CHECK(grib_set_long(h,"gridDefinition",255),0);
```

grib_ls

grib_ls test.grib1 tigge_pf_ecmwf.grib2

```
test.grib1
centre    param      levelType   level   date       dataType   short_name   valuesCount   gridType      packingType
ecmf      130.128    ml          1       20070323   an         T            6114          reduced_gg    grid_simple
ecmf      130.128    pl          1000    20070323   an         T            6114          reduced_gg    grid_simple
ecmf      167.128    sfc         0       20070323   an         2T           6114          reduced_gg    grid_simple
ecmf      229.140    sfc         0       20070323   an         SWE          311362        reduced_ll    grid_simple
ecmf      130.128    ml          1       20070323   an         T            8192          regular_gg    grid_simple
ecmf      130.128    pl          1000    20070323   an         T            8192          regular_gg    grid_simple
ecmf      167.128    sfc         0       20070323   an         2T           8192          regular_gg    grid_simple
ecmf      167.128    sfc         0       20070323   an         2T           496           regular_ll    grid_simple
ecmf      167.128    sfc         0       20070323   an         2T           496           regular_ll    grid_simple
9 of 9 grib messages in test.grib1

tigge_pf_ecmwf.grib2
centre    date       dataType    gridType     step    levelType   lev     short_name   packingType    valuesCount
ecmf      20070122   pf          regular_ll   96      sfc         10      10u          grid_simple    684
ecmf      20070122   pf          regular_ll   96      sfc         10      10v          grid_simple    684
ecmf      20070122   pf          regular_ll   96      sfc         0       cape         grid_simple    684
ecmf      20070122   pf          regular_ll   96      pl          925     gh           grid_simple    684
...
ecmf      20070122   pf          regular_ll   96      pl          925     u            grid_simple    684
ecmf      20070122   pf          regular_ll   96      pv          2       u            grid_simple    684
ecmf      20060619   pf          regular_ll   96      sfc         0       u            grid_simple    684
ecmf      20070122   pf          regular_ll   96      pl          925     v            grid_simple    684
ecmf      20070122   pf          regular_ll   96      pv          2       v            grid_simple    684
38 of 38 grib messages in tigge_pf_ecmwf.grib2

47 of 47 total grib messages in 2 files
```

grib_ls (mars keys)

grib_ls -m test.grib1

```
test.grib1
param      levtype    levelist   date       time   startStep   endStep   step   class   type   stream   expver
130.128    ml         1          20070323   1200   0           0         0      od      an     oper     0001
130.128    pl         1000       20070323   1200   0           0         0      od      an     oper     0001
167.128    sfc        0          20070323   1200   0           0         0      od      an     oper     0001
229.140    sfc        0          20070323   1200   0           0         0      od      an     wave     0001
130.128    ml         1          20070323   1200   0           0         0      od      an     oper     0001
130.128    pl         1000       20070323   1200   0           0         0      od      an     oper     0001
167.128    sfc        0          20070323   1200   0           0         0      od      an     oper     0001
167.128    sfc        0          20070323   1200   0           0         0      od      an     oper     0001
167.128    sfc        0          20070323   1200   0           0         0      od      an     oper     0001
9 of 9 grib messages in test.grib1

9 of 9 total grib messages in 1 files
```

options

A set of options common to all the tools is provided.

- -p key[:{s/d/l}],key[:{s/d/l}],...

Declaration of keys to print. For each key a string (key:s) or a double (key:d) or a long (key:l) value is printed. The default type is string.

- -w key[:{s/d/l}]{=/!=}value,key[:{s/d/l}]{=/!=}value,...

Where clause. Grib messages are processed only if they match all the key/value constraints. A valid constraint is of type key=value or key!=value. For each key a string (key:s) or a double (key:d) or a long (key:l) type can be specified. The default type is string.

grib_ls (-p option)

grib_ls -p count,centre test.grib1

```
test.grib1
count       centre
1           ecmf
2           ecmf
3           ecmf
4           ecmf
5           ecmf
6           ecmf
7           ecmf
8           ecmf
9           ecmf
9 of 9 grib messages in test.grib1

9 of 9 total grib messages in 1 files
```

grib_ls -p count,centre:l test.grib1

```
test.grib1
count       centre
1           98
2           98
3           98
4           98
5           98
6           98
7           98
8           98
9           98
9 of 9 grib messages in test.grib1

9 of 9 total grib messages in 1 files
```

grib_ls (-w option)

grib_ls -w levelType=pl test.grib1

```
test.grib1
edition   centre  param    levelType   level  date      dataType  short_name  valuesCount  gridType     packingType
1         ecmf    130.128  pl          1000   20070323  an        T           6114         reduced_gg   grid_simple
1         ecmf    130.128  pl          1000   20070323  an        T           8192         regular_gg   grid_simple
2 of 9 grib messages in test.grib1

2 of 9 total grib messages in 1 files
```

grib_ls -w levelType!=pl test.grib1

```
test.grib1
edition  centre  param     levelType  level  date      dataType  short_name  valuesCount  gridType     packingType
1        ecmf    130.128   ml         1      20070323  an        T           6114         reduced_gg   grid_simple
1        ecmf    167.128   sfc        0      20070323  an        2T          6114         reduced_gg   grid_simple
1        ecmf    229.140   sfc        0      20070323  an        SWH         313362       reduced_ll   grid_simple
1        ecmf    130.128   ml         1      20070323  an        T           8192         regular_gg   grid_simple
1        ecmf    167.128   sfc        0      20070323  an        2T          8192         regular_gg   grid_simple
1        ecmf    167.128   sfc        0      20070323  an        2T          496          regular_ll   grid_simple
1        ecmf    167.128   sfc        0      20070323  an        2T          496          regular_ll   grid_simple
7 of 9 grib messages in test.grib1

7 of 9 total grib messages in 1 files
```

grib_copy (-w option)

grib_copy -w levelType=pl test.grib1 pl.grib1

```
test.grib1
edition  centre   param    levelType level   date     dataType  short_name  valuesCount  gridType     packingType
1        ecmf     130.128  pl        1000    20070323 an        T           6114         reduced_gg   grid_simple
1        ecmf     130.128  pl        1000    20070323 an        T           8192         regular_gg   grid_simple
2 of 9 grib messages in test.grib1

2 of 9 total grib messages in 1 files
```

grib_copy -w levelType!=pl test.grib1 not_pl.grib1

```
test.grib1
edition  centre  param      levelType  level  date      dataType  short_name  valuesCount  gridType     packingType
1        ecmf    130.128    ml         1      20070323  an        T           6114         reduced_gg   grid_simple
1        ecmf    167.128    sfc        0      20070323  an        2T          6114         reduced_gg   grid_simple
1        ecmf    229.140    sfc        0      20070323  an        SWH         313362       reduced_ll   grid_simple
1        ecmf    130.128    ml         1      20070323  an        T           8192         regular_gg   grid_simple
1        ecmf    167.128    sfc        0      20070323  an        2T          8192         regular_gg   grid_simple
1        ecmf    167.128    sfc        0      20070323  an        2T          496          regular_ll   grid_simple
1        ecmf    167.128    sfc        0      20070323  an        2T          496          regular_ll   grid_simple
7 of 9 grib messages in test.grib1

7 of 9 total grib messages in 1 files
```

grib_set (-w option)

grib_set -s date=20070523 -w levelType=pl test.grib1 pl.grib1

grib_ls pl.grib1

```
pl.grib1
edition centre   param    levelType  level  date      dataType  short_name  valuesCount  gridType     packingType
1       ecmf     130.128  ml         1      20070323  an        T           6114         reduced_gg   grid_simple
1       ecmf     130.128  pl         1000   20070522  an        T           6114         reduced_gg   grid_simple
1       ecmf     167.128  sfc        0      20070323  an        2T          6114         reduced_gg   grid_simple
1       ecmf     229.140  sfc        0      20070323  an        SWH         313362       reduced_ll   grid_simple
1       ecmf     130.128  ml         1      20070323  an        T           8192         regular_gg   grid_simple
1       ecmf     130.128  pl         1000   20070522  an        T           8192         regular_gg   grid_simple
1       ecmf     167.128  sfc        0      20070323  an        2T          8192         regular_gg   grid_simple
1       ecmf     167.128  sfc        0      20070323  an        2T          496          regular_ll   grid_simple
1       ecmf     167.128  sfc        0      20070323  an        2T          496          regular_ll   grid_simple
9 of 9 grib messages in pl.grib1

9 of 9 total grib messages in 1 files
```

grib_set (changing packingType)

grib_set -s packingType=grid_jpeg test.grib2 test_jpeg.grib2

grib_ls test.grib2 test_jpeg.grib2

```
test.grib2
edition centre   date      dataType gridType    levelType  lev  packingType  valuesCount
2       ecmf     20070323  af       reduced_gg  sfc        0    grid_simple  6114
1 of 1 grib messages in test.grib2

test_jpeg.grib2
edition centre   date      dataType gridType    levelType  lev  packingType  valuesCount
2       ecmf     20070323  af       reduced_gg  sfc        0    grid_jpeg    6114
1 of 1 grib messages in test_jpeg.grib2
2 of 2 total grib messages in 2 files
```

grib_set (setting MISSING)

grib_set -s scaleFactorOfSecondFixedSurface=missing test.grib2 test_m.grib2

grib_ls test.grib2 test_m.grib2

```
test.grib2
scaleFactorOfSecondFixedSurface
0
1 of 1 grib messages in test.grib2

test_m.grib2
scaleFactorOfSecondFixedSurface
MISSING
1 of 1 grib messages in test_m.grib2

2 of 2 total grib messages in 2 files
```

grib_set (changing precision)

grib_set -s setDecimalPrecision=1 test.grib2 test_prec1.grib2

grib_compare test.grib2 test_prec1.grib2

```
[referenceValue] double values are different: [21025.5] and [2102.55], diff:
     18922.9
[decimalScaleFactor] long  values are different: [2] and [1]
[numberOfBitsContainingEachPackedValue] long  values are different: [14] and [11]
[values] 6114 values, max absolute difference: 0.0499951
```

GRIB Tools: grib_convert

converts grib messages, applying the rules from a conversion_rules file, e.g:

```
editionNumber = 2;
if( indicatorOfParameter == 11 && indicatorOfTypeOfLevel == 105)
{
   productDefinitionTemplateNumber = 1;
   typeOfFirstFixedSurface         = 103;
   scaleFactorOfFirstFixedSurface  = 0;
   scaledValueOfFirstFixedSurface  = 2;
}
```

grib_filter

is similar to grib_convert.

It supports more complex rules and is used to split a single file into several files with a write rule:

write "[centre]_[date]_[dataType]_[levelType].grib[editionNumber]";

It will produce the following set of files from a single file:

ecmf_20060619_pf_sfc.grib2

ecmf_20060630_pf_sfc.grib2

ecmf_20070122_pf_pl.grib2

ecmf_20070122_pf_pt.grib2

ecmf_20070122_pf_pv.grib2

ecmf_20070122_pf_sfc.grib2

grib_get_data

prints a latitude/longitude/value list.

It takes into account the bitmap.

The -m option provides a way to define a missing value: grib_get_data -m 9999:missing file.grib2 prints "missing" every time a 9999 is found.

Future developments

- Replace gribex in the operational applications.
- Implement thread safety in the Fortran interface. The C interface is already thread safe.
- Implement memory management. The library is already designed to change the memory allocation function with user defined functions.
- Modify the iterators to deal with all the scanning modes and to change from one scanning mode to another.
- Provide a function to get the nearest neighbours of a latitude/longitude point.
- Implement a GUI based on the GRIB API to inspect and modify grib files easily.
- Work on the documentation, providing more examples.
- Develop a Fortran 90 interface. A Fortran 77 interface is already available and can also be linked to Fortran 90 programs.
- Release a perl interface: it is already implemented and tested and will be released in the next version.
- Design and implement more interfaces (matlab, java, python ?)
- Distribute the library with the LGPL licence.

T. Lorenzen asked a related question; he believed that models running on the new HPCF in two years' time would have to produce output in GRIB 2. When did ECMWF believe that users would have to be ready for this? *U. Modigliani replied that he couldn't give an exact date but users would get plenty of notice; GRIB 1 would certainly be acceptable for more than another two years for current products only.*

ECaccess: status and plans – *Laurent Gougeon*

ECaccess developments

- Initially used to implement a secure portal to access ECMWF archiving and computing facilities
  - Providing file and job management in batch or interactive mode
  - Supporting secure interactive sessions on ECMWF platforms
- Implementation of a secure portal to MS
  - Allowing secure file transfers between ECMWF and remote sites
  - ECaccess network with the ECaccess Gateways
- Set of configurable Java components for various ECMWF projects
  - Services and libraries for security, file transfer, data management, job submission and control, monitoring ...
- Common CVS repository
  - More than 220,000 lines of code (1,500 Java classes) using 100 external libraries (Open source API)

Uses of ECaccess

- Secure portal to ECMWF (ECaccess)
  - Via the Internet or RMDCN, ECaccess provides file and job management in batch or interactive mode and allows secure file transfers between ECMWF and remote sites
- ECMWF Product Dissemination System (ECPDS)
  - Provides a data transmission service which allows MSs to specify which ECMWF product to deliver, on which target systems, using which networks (RMDCN, Leased Lines, Internet/ECaccess, LAN)
- MS Time-Critical Applications (MSTCA)
  - MS job submissions triggered by SMS suites and monitored by ECMWF operators
- Projects using ECaccess components
  - PKI (Web Services Gateway), Strong Authentication (Web Help-desk and Self-desk), Ticketing Server ...

ECaccess statistics

- ECaccess
  - More than 2,000 users recorded in ECaccess
  - 40,000 jobs/month on average on ecgate and hpce
  - 395,000 transfers/month on average with ectrans
  - 140,000 transfers/month on average with ftp
- ECPDS (transfers more than 7,100 GB/month)
  - Internet (Local Gateways): 235,000 files/month (4,500 GB/month)
  - LAN: 15000 files/month (1,600 GB/month)
  - RMDCN: 320,000 files/month (675 GB/month)
  - Leased Lines: 41,500 files/month (280 GB/month)
  - Internet (Remote Gateways): 4,500 files/month (18 GB/month)
- MSTCA
  - More than 600 jobs with 43 daily events and 2 weekly events

ECaccess network update

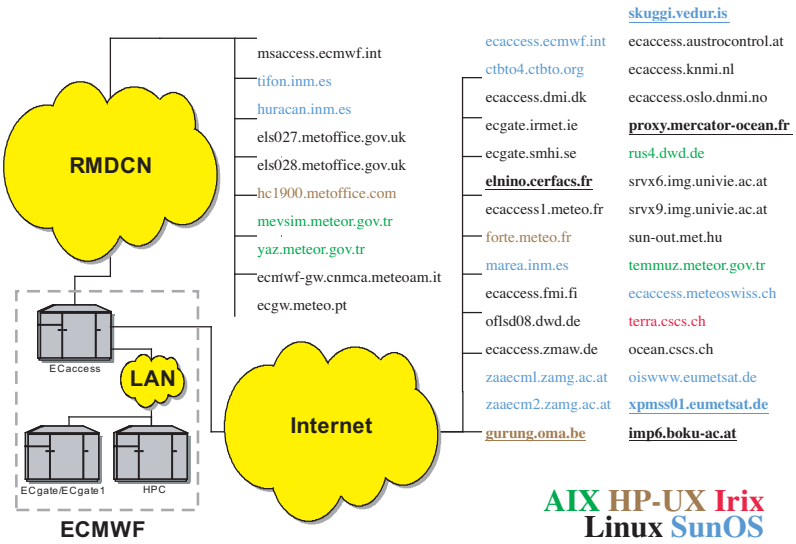

Upgrade to ECaccess v3.1.0

- The upgrade has been mandatory to support the new strong authentication mechanism based on ActivID.
  - The Telnet, ssh plugins and eccert command were restricting passwords to 6 digits (avoiding unnecessary network traffic for invalid SecurID passcodes).
- Almost all ECaccess Gateways have been upgraded.
- There are some issues regarding the upgrade of old ECaccess Gateways running with Java 1.3. The UKMO has additional problems, due to a special set-up on their side, originally implemented to fix bad performance over their Leased Line.
- Remote sessions to assist MS in their upgrade
- Questions regarding “eccert -expire” from INM/DWD:
  - Validity for each command and exit code?
- 5 Gateways left to upgrade!

Progress on ECaccess New Generation plans

- ECaccess Virtual File System (ECVFS)
  - ECVFS at the same level as ECHOME/ECSCRATCH
  - Transfers possible between different domains and different ECaccess Gateways (through ectrans associations)
- Progress on a new Web interface (Web 2.0)
  - More responsive and more user friendly (e.g. cut and paste between ECaccess domains)
  - Based on AJAX technologies
- Implementation of Web Services
  - Services provided by the ECaccess Gateway now available through SOAP requests (a step towards the integration of Globus)
  - Based on AXIS (SOAP implementation used by Globus)
  - Could be used for the new release of the ectools
- VNC Applet improvements
  - Support for slow/fast networks

- ECFS improvements
  - Support for meta-characters in files names and support for multiple deletions at once (e.g. mdel *)
- Integration with the new ECMWF PKI
  - Better management of the Gateway certificates due to expire
  - Will allow MS administrators to create certificates for temporary users (e.g. for the ectools)
- Additional ectrans module
  - The ftps protocol is now supported (ftp over ssl)
- Release of the ECMARS plug-in
  - Originally developed for DWD, this plug-in will be part of the standard ECaccess distribution and will allow the MARS archive to be queried by the ectools
- Added support for the new hpcf

Problems already/almost fixed

- Large increase in usage has triggered new problems
- ECgate (ECaccess Resources Provider)
  - OS overload has sometimes caused delays in job sub-missions (e.g. for INM)
  - Impossible to upgrade to Java 1.5 (system upgrade required)
- ECaccess Job Scheduler
  - During the migration to the new MSTCA we have fixed a number of issues regarding the integration with LL
- Gateway certificate management
  - A few sites (e.g. KNMI, DWD, CTBTO) have been affected by the sudden expiration of their Gateway certificate without warning, messages will be sent in advance
  - Gateway certificates have been extended to give time for their replacement without affecting users
  - Gateway certificates from the new ECMWF PKI will soon be distributed to MS with a new duration of 4 years instead of 2 years

Remaining issues

- ECPDS
  - Problems during dissemination peeks on the Internet link causing unnecessary alerts for the operators (unnoticeable for the users, as the transfers are automatically restarted almost instantly)
  - Still investigating: network congestion, rooter or other security device, bug in RMI implementation, ...
- Database
  - No redundant system to take over the service
- Timeouts
  - Transfers with ectrans: some users at MF are experiencing interruptions/retry with their transfers
  - Telnet and ssh: the standard firewall timeout of 10 minutes applies for the ECaccess data channel on port 9003 (e.g. EUMETSAT)

What next?

- We will work hard to fix the remaining issues!
- ECaccess Resource Provider will be moved to the new MS server
  - Support for Java 1.5 (the application has been already tested successfully on ECgate1)
  - A more powerful system which should solve the availability issues for ECaccess
- ECaccess v3.2.0 scheduled later this year
  - New features (e.g. ECaccess VFS and the new Web interface) will be gradually deployed locally at ECMWF and, if validated, incorporated in the new distribution
  - Java requirements will be the same as for the ECaccess v3.1.0 Gateways!
- Continue progressing on last year's plans: finalize the Globus integration, new ectools, load balancing ...

ECaccess links

- **http://www.ecmwf.int/services/ecaccess/**
  - Download the ECaccess packages, the user's manual (on-line & PDF) and the administrator's manual (PDF)
  - Registration centre (SecurID or ActivID card required)
- **http://ecaccess.ecmwf.int/sshvnc/**
  - Interactive session (vnc over ssh)
- **https://ecaccess.ecmwf.int:9443/**
  - Dissemination monitoring
- **https://ecaccess.ecmwf.int:8443/**
  - ActivID Web helpdesk
- For access through the RMDCN instead of the Internet, just replace ecaccess by msaccess

Member State Time Critical Applications – *Dominique Lucas*

Background

- Framework for time-critical applications:
  - Demand for more robust environment for job submission, with monitoring facilities
  - Requests from some Member State users to set up more developed time critical systems at ECMWF
  - No satisfactory solution available
  - Discussed at a TAC meeting and approved by Council in 2004
  - 3 options considered for the framework:
    - Option 1: Simple time critical jobs
    - Option 2: SMS suites, managed by users, monitored by ECMWF
    - Option 3: SMS suites, managed and monitored by ECMWF

Simple Time Critical jobs

- Service based on ECaccess - reminder
  - Enhancement of batch service under ECaccess
  - Notion of events/notifications added
  - Users subscribe jobs to the events
  - ECMWF operational suite sends a notification to an event
  - This notification will trigger the user jobs
  - These jobs are visible to our operators
- Introduction of the service in November 2006
- Reminders have been sent to users of the old system
- Old system closed on 9 May 2007

Simple Time Critical jobs – ECtools

- ecels – list events available to user:

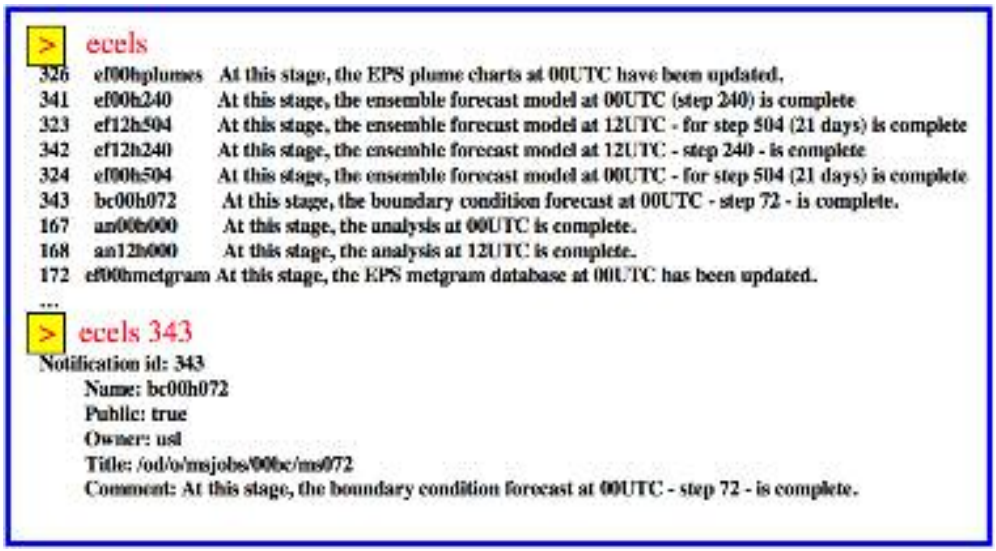

- ecjput – ecjls

```
> ecjput -help
Syntax: JPUT ECaccess-queue local-script [args ...]
...
-ni - notifications ids (list separated by ';' or ',')
-ro - renew subscription off (default is on)
-oo - one script to one notification off (default is on)
-mr - send mail when the execution/transfer retries
-mp - man page content (comment for the operators)
-rc - define the number of job retries (default is 0)
-rf - define the frequency of job retries in seconds (default is 600)
...
> ecjput ecgate sms.cmd -nd -ni 343 -mp "nothing to be done" -rc 1 -rf 300
35853
> ecjls 35853
Jobid: 35853
Location: ecgate@ecgate.ecmwf.int
Notification(s): bc00h072 (343)
Schedule: May 31 20:06
Try number: 0/2
Status: STDBY
```

Simple Time Critical jobs – today

- 94 different users of this service
- More than 600 jobs: a dozen on hpce, the rest on ecgate
  - Mainly simple, straightforward jobs
  - A few jobs with dependencies:
    - Jobs with multiple steps (Loadleveler feature)
    - Jobs with built-in dependencies, maintained by the user: jobs submitting other jobs ...
- Monitoring:
  - Operators monitor the jobs and notify User Support of any failure.
  - User Support contacts the users.
  - Errors mainly occur in the installation period.
  - Users are given ecommendations to increase the reliability of their jobs.

Simple Time Critical jobs – issues

- Common problems in migration:
  - “set -e” stopped some jobs on harmless errors
  - Loadleveler directives/no directives in job
  - Jobs submitted to wrong event - data not available
- ECaccess (specific) job statuses:
  - STDBY: job with subscription to an event, waiting for notification
  - INIT: job initialised by ECaccess, being submitted to batch system on target host, e.g. Loadleveler on ecgate.
  - WAIT: job is queued in batch system on target host
  - EXEC: job is running
  - DONE: Job is complete and finished successfully
  - STOP: Job is complete but failed.

- Increased complexity in user jobs:
  - Multi-step jobs:
    - Operators only have visibility of one ECaccess job
    - If the job (any of the steps) fails, operators may restart the job, including all the job steps => increased use of resources, "waste of time"
    - Check previous runs of job step within the job ...
  - Jobs submitting other jobs:
    - Operators only have visibility of one ECaccess job, the first user job. Any failure in subsequent jobs will not be visible.
    - Use job steps ...

Simple Time Critical jobs – ecgate

- Peak times on ecgate:
  - in the morning and evening
  - may slow down user Time Critical jobs
  - New ecgate should absorb more "traffic".
  - We may look at prioritising the Time Critical jobs.
- Long running jobs:
  - Access to VAREPS and Monthly forecast data may take some time.
  - Single jobs will run faster on the new ecgate.
  - Users can divide their work into separate jobs.

SMS suites

- Suitable for more complex Time Critical activity
- To be discussed with User Support; initial test suites can be developed.
- Service to be requested officially by TAC representatives via our Head of Operations Department
- SMS should be used ...
- Some guidelines should be followed. These guidelines are based on the experience acquired at ECMWF in running our operational activity.

Environment

- Initial tests can be done with normal UID.
- Time critical activity will be implemented with a special UID.
- Access to the 2 HPCF clusters will be opened.
- Access to privileged classes/queues given.
- Access to special file system will be given, with quotas!
- SMS:
  - To be run on ecgate
  - Generic script available to start SMS
  - Controlled assignment of SMS program numbers
  - Generic Log server available on HPCF system
  - Generic submission script on ecgate
  - Generic script to kill a job from xcdp
  - SMS access protections (sms.lists file)
- Monitoring will be done by the operators.

Recommendations

- A document detailing the set up of a suite in SMS has been written and will be published on our web site shortly.
- Use variables in SMS to change the configuration of your suite, the profile of a job ...
- Document all the tasks in your suite, with clear instructions to the operators.
- Run the critical path in the safest environment - on HPCF.
- Do not access the DHS (mars or ECFS) in your critical path.
- Do not use cross-mounted file systems. Use local file systems.
- Keep a tight control on disk usage, including data and SMS related files.
- The triggering of your suite can happen with a simple time critical job or with a time trigger in your suite.
- The cycling of your suite should happen with a time trigger.
- Data transfers:
  - With external sites: use ectrans
  - Between systems at ECMWF: use ecrcp
- Backup procedures should be implemented to use the second HPCF cluster.
- Changes to the suite should first be done in test mode.
- A mailing list will be set up to exchange relevant information.

Issues

- Data transfers: ectrans
  - Ectrans (being asynchronous) is not monitored by operators.
  - A dissemination-like system could be implemented.
  - Use synchronous ectrans transfers
- Ecgate: single point of failure
  - SMS runs on ecgate ...
  - Some tasks in the critical path may run on ecgate.
  - The Linux cluster should provide higher resilience.

V. Gislason asked whether the new SMS document which D. Lucas had mentioned referred to SMS in general or only to time critical SMS. *D. Lucas replied that the new document was specific to the setting up of suites in a time critical environment; general SMS documentation remains available on the ECMWF public website.*

New Incident Management System – *Paul Dando*

Background

- ECMWF has practised Incident Management for many years – we have to!
- Use of software solutions has improved efficiency.
- Formal incident management has matured only recently.
- Consideration of such an approach can help to identify the strengths and weaknesses of our processes.
- Current processes result from:
  - the support requirements of our customers
  - the currently used software solutions
  - not everyone using the current software solutions
- This results in a fragmented approach.

History

- Incident management tools (Operators and Calldesk)
  - 1980 – Paper based system
  - 1992 – Repgen: Empress 4GL forms system
  - 1999 – Repgen-2: Web-based system
  - 2002 – Calldesk shared IMAP e-mail knowledge base added
  - 2003 – Analysts' Repgen: Analysts' web-based interface
- Incident management tools (Analysts)
  - 1980 – Face-to-Face and provider facilities (EQUANT, JANET, IBM)
  - 1992 – E-mail to log – Change notices
  - 2001 – Web shared IMAP e-mail
  - 2003 – Analyst Repgen – read/reply

The need for a new tool

- To allow all support persons to record all incidents in one system
- To allow the recording of incidents without unnecessary overheads
- To allow recorded incidents to be used as a knowledge base by all support staff
- To facilitate the formal assignment of responsibility for and documentation of the resolution of recorded incidents
- To provide a means of analysing our service quality
- To provide a means of creating server availability statistics
- To provide a means of producing customised incident reports used by the Calldesk

Incident Management project: scope

- Included
  - Incident management
  - Knowledge management relating to incidents
  - Users: Computer Division (and Meteorological Division)
- Excluded
  - Change management
  - Configuration management
  - Problem management
  - Release management
  - Asset management
  - Formal Service Level Agreements

Core project team

The core project team consisted of representatives from:

- Shift staff ("24/7" Operators)
- Calldesk
- Analysts
- User Support

Selection process

- Document current processes
- Define requirements
- Produce a Request for Proposal (RFP)
- Undertake initial market survey
  - 24 products identified as having the potential to meet our requirements
  - Copies of the RFP were sent to vendors
- Six responses to the RFP were received from vendors:
  - LiveTime, ServiceDesk Plus, TOPdesk and Monitor 24-7
  - Two vendors offered Remedy - but only it supports MS Internet Explorer
- Two RFPs were evaluated internally
  - Current in-house solution (Repgen) and Request Tracker

Detailed evaluation

- Detailed evaluations were undertaken to test:
  - Ease of customisation
  - Usability by Shift Staff
  - Usability by Calldesk
  - Usability by User Support and Analysts
- Tested: ServiceDesk Plus, TOPdesk, Monitor 24-7 and LiveTime
  - Both Repgen (known) and Request Tracker required significant in-house development
- Serious shortcomings were found in all products
- Conclusion: none of the products evaluated was suitable for deployment at ECMWF

Review of other products

- Several of the companies that had not responded to the RFP were contacted:
  - To establish their reasons for not responding
  - To check whether their product could meet our requirements
- Our evaluation was based on what we had learnt from the detailed evaluation of the other products and two potential candidates were identified:
  - Footprints by Unipress (now Numara) – Linux-based
  - Novo HelpDesk by Novo – Windows-based
- It was decided to evaluate Footprints further.

Footprints – general impressions

- Easy to install, maintain, configure and customise
- Very intuitive to use - very little training needed
- ITIL compatible - PinkVerify(tm) certification

- Provides most of the features requested in the RFP, some shortcomings have been identified
- Accessible for code changes
- Uses software with which we are familiar: code written in Perl and CGI; uses PostgreSQL
- Provides comprehensive features for managing incidents using e-mail
- Well-supported by the vendor
- Chosen!

Footprints implementation

- Footprints has been purchased and installed at ECMWF
- Initial configuration and testing has taken place
- First trial involving analysts has been completed: no major problems identified
- Currently reviewing the configuration and settings
- Some technical issues with the Operator environment need to be resolved
- Plan to implement operationally in mid-late June (TBC)
  - Initially only for Operators, Calldesk and analysts
  - Access for internal and Member State users to be reviewed later

D. Birman asked whether any organisation of responsibilities had been defined around the new tool, for instance, who was responsible for the progress tracking, and asked P. Dando to describe the 'life' of an incident. *P. Dando replied that the rôles hadn't been finalised but the Call Desk usually took responsibility for the main tracking. Guidelines state that the first person to notice a problem should be the one to open the incident report. The usual 'life' of an incident would be: problem observed on the supercomputer, probably by the shift staff; details of the incident entered into Footprints; if the area of responsibility is obvious, the problem may be assigned to the appropriate analyst immediately. This is not current working practice, which is that the Call Desk deals with the assignment of all incidents. The designated analyst then works on the problem, assigning it sideways, if necessary, and logs its resolution, once complete. The Call Desk's rôle will probably be to review and 'chase' outstanding incidents, though this procedure has not yet been finalised.*

R. Rudsar asked how the knowledge base built up from past experience would be used. P. Dando replied that the Footprints database had an excellent search facility for discovering previous incidences of similar problems. R. Rudsar noted that that in Norway they had experienced some difficulty in relocating specific past problems and had introduced a one-word keyword for each problem. *P. Dando replied that a sophisticated search facility had been a very high priority, when assessing potential replacements for Repgen. The situation may have to be re-assessed, once the database has grown and searches take increasingly longer. I. Weger added that Footprints' search facility searched on both metadata and full text.*

D. Birman explained that Météo France had selected Remedy to be used by users with little knowledge of computers. Later, different levels of display were added for more knowledgeable users. (For more information on Remedy see Annex 3).

P. Dando added that the ECMWF Footprints facility also had different displays for different users. P. Halton suggested that a pull down menu of problem definitions could be created to aid problem categorisation. I. Weger replied that this had already been done, using the categorisations from Repgen and building on them.

Usage of HPCF resources by Member States – *Dr Umberto Modigliani*

HPCF resources available to Member States

- Council guidelines on the distribution of computer resources establish that:
  - "at least 25% of the supercomputer system shall be made available to the Member States, as well as sufficient resources for data storage ..."
  - A maximum of 10% of the computer resources available to the Member States may be allocated for "special projects"
- 35% of the supercomputer resources is allocated equally among the Member States and 65% allocated proportionally to their financial contribution
- Special Projects:
  - "are of a scientific or technical nature, undertaken by one or more Member States and likely to be of interest to the general scientific and/or technical community"
  - Requests for such projects have to be submitted annually, as part of Member State estimates of their future use of the Centre's computing resources.
  - These projects are then considered by Council at its autumn session, prior to the year in which the Special Project will take place.
  - 20% of the available Special Project resources are set aside for late applications and can be granted at any time of the year.

Usage of HPCF resources by MS users

Research activities

- Research in numerical weather forecasting (NWP) by National Meteorological Services (NMS)
  - Aladin
  - COSMO/Lokal Modell
  - HIRLAM
  - Unified Model
  - Other (IFS, Arpege, GME, MM5, WRF, RAMS, ROMS, etc)
- Research in NWP by other users (universities, etc.) using models such as ECHAM, ECHAM5-HAM, MesoNH, MM5, MPI-OM, OPA, CLM, HIRHAM, RACMO2, RegCM, Speedy, MICOM, MOZART, TM5, etc.

Special Projects

- 89 Special Projects in 2007, of which 75 are continuations of projects started previously
- Submitted by researchers in 13 Member States + JRC and ICTP
- Researchers from Co-operating States can have an active role in a Special Project and, therefore, have access to HPCF resources
  - SPFRCOUP: "Investigation of coupling the ALADIN and AROME models to boundary conditions from ECMWF and ERA model data", includes researchers from: Hungary, Slovenia, Croatia, Romania, Czech Republic, Morocco, Slovakia
- There is a very large variety of project interests, covering NWP, ocean modelling, climate research, post-processing of data, etc.
- A list of all Special Projects, describing their activities, is available at:

  www.ecmwf.int/about/special_projects/
- Guidelines on how to apply are available at:

  www.ecmwf.int/about/computer_access_registration/Special_Projects.html
- See article "Applying for resources for a Special Project" in ECMWF Newsletter No. 110 Winter 2006/07 available at:

  www.ecmwf.int/publications/newsletters/pdf/110_rev.pdf

Time Critical applications

- Framework for Member State time critical applications
- COSMO-LEPS 16 LM at 10km/40ML up to 132h once a day at 12 UTC
- Norwegian TEPS/LAMEPS
- 20+1 TEPS members at T399L62 up to 72h once a day at 12 UTC
- MOTHS: EPS based on UM as the UKMO contribution to TIGGE
- 23+1 UM members at N144L38 up to 360h twice a day at 0/12 UTC
- COSMO-MED/COSMO-ITA using specific 3D-Var analysis and ECMWF BC for the Italian Met Service
- 3D-Var assimilation at 14 km every 3 hours
- COSMO-MED at 7km up to 78h twice a day at 0/12 UTC
- COSMO-TA at 2.8 km up to 36h once a day at 0 UTC
- Aladin LAEF for the Austrian Met Service
- 18 members, 18km/37ML twice a day at 0/12 UTC (option-1)
- EUROSIP multi-model seasonal forecasts: ECMWF, Meteo-France and UKMO

Optional Projects

- Boundary Conditions "Optional Project"
- Boundary value forecasts from the 4 main synoptic hours

(0/6/12/18 UTC) up to 90h

- Using 4D-Var since March 2006
- Supported by 15 Member States (all but France, Germany, UK) and 7 Co-operating States
- Using ~16% of the participating countries' HPCF resources in 2007

Summary

Country	Research	Special Project	Time-critical	Optional Project
Belgium	MUMM (North Sea modelling)			Yes
Denmark	HIRLAM	1		Yes
Germany	GME, GME-EPS, LME, prepIFS	29	COSMO-LEPS	No
Spain	HIRLAM, Aladin, RCA seas, traj	3		Yes
France	Aladin, Olive, Mocage climate, ocean modelling	4	EUROSIP	No
Greece	LM, traj		COSMO-LEPS	Yes
Ireland	HIRLAM, climate/ocean modelling	2		Yes
Italy	HRM, LM	11	COSMO-LEPS COSMO-MED/ COSMO-ITA	Yes
Luxembourg				Yes
Netherlands	HIRLAM, CTM, IFS, EC-EARTH?	10		Yes
Norway	HIRLAM, IFS (+CTM)	5	TEPS/LAMEPS	Yes
Austria	Aladin	8	Aladin LAEF	Yes
Portugal		1		Yes

Table continued on next page

Country	Research	Special Project	Time-critical	Optional Project
Switzerland	LM reforecasts, EFI	1	COSMO-LEPS	Yes
Finland	HIRLAM (RCR), MM5 (+LAPS)	1		Yes
Sweden	HIRLAM, EC-EARTH?			Yes
Turkey	MM5			Yes
United Kingdom	Atmos, coupled ocean-atmos	10	EUROSIP MOTHS	No

Usage statistics

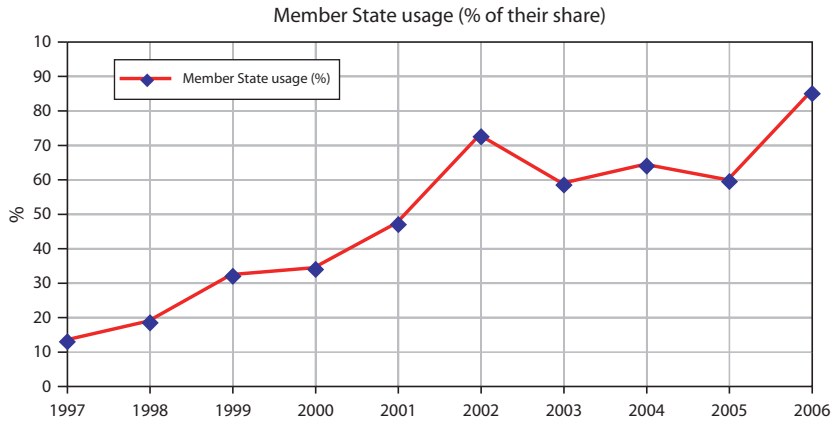

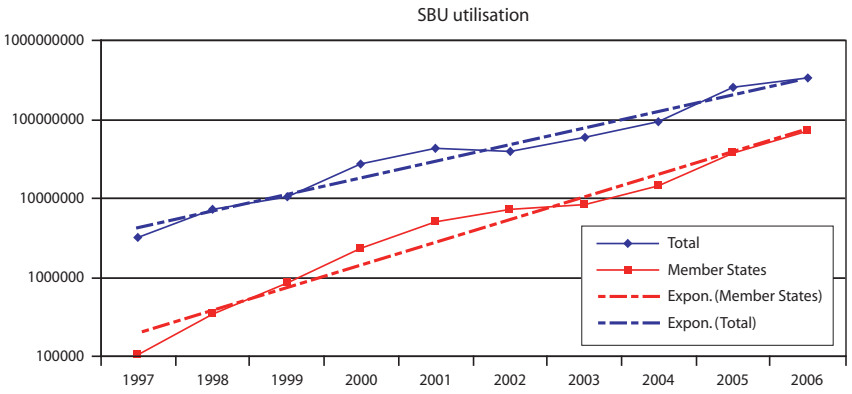

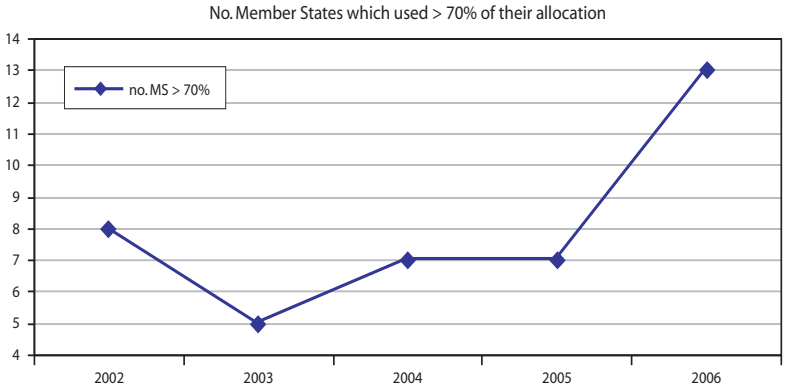

Discussions – *Umberto Modigliani*

ecgate

- More resources requested:
  - Disk quota
  - Number of jobs running per user
  - Faster processors

U. Modigliani acknowledged the importance of the ecgate server to Member State users. In response to some representatives' requests for various additional resources he noted that R. Fisker's presentation would have given more information on current plans to replace the complete server. In the last few months various unpopular but unavoidable restrictions have been placed on users' working practices (e.g. limiting the number of jobs that an individual user can run at any one time) in order to mitigate the effects of the overload and will be removed once more resources become available.

R. Rudsar observed that some of her users had been advised to run MARS jobs on HPCE. Would this advice remain valid, once ecgate had been upgraded? U. Modigliani replied that it would to some extent depend on the purpose of the data being retrieved. If they were required for an application to run on the HPC, then they could be retrieved there. Otherwise, MARS jobs should probably return to ecgate.

G. Wotawa asked whether there would be more memory on the new ecgate. U. Modigliani replied that there would be more memory on ecgate but reminded representatives that the default memory addressing mode was 32-bit. The aim is eventually to support a 64-bit environment and, in this case, it would be possible to exceed the normal 32-bit system 2 Gigabyte limit. R.Fisker noted that ECMWF is still considering the replacement of ecgate by a Linux cluster, each node of which would be likely to have a maximum of approx. 8 GB, so there would only be 7 GB available for individual jobs. Guidelines would be produced, in the light of experience with the new system. In the meantime, users with specific, unusual requirements should contact User Support for advice.

MARS

Speed of interpolation from regular or rotated grid to lat/long, particularly for wind components

U. Modigliani explained that interpolation within MARS is a serial program and is therefore slower than interpolation during product generation, which is parallelized over up to three nodes. T. Lorenzen asked whether it would be possible to parallelize interpolation in MARS. U. Modigliani replied that this would not be easy but could be investigated, if there was some specific requirement for wind components. The current interpolation library is being rewritten, which will benefit MARS; once this is complete, parallelization might be investigated but this would not happen in the near future.

D. Lucas observed that users could submit their MARS jobs to HPCF in MPMD mode using LoadLeveler, which allows several MARS jobs to be run in parallel, producing an environment similar to that of product generation.

Access to MARSLIB at the same level as EMOSLIB

U. Modigliani stated that this would need to be discussed internally. Personally, he preferred not to proliferate the number of libraries available to users and proposed that those routines in MARSLIB which were of particular use to users could be added to EMOSLIB. He pointed out that with the introduction of GRIB-API there is also a plan to split EMOSLIB into its separate components and this would provide an opportunity to creating a stand-alone MARSLIB.

Why are there no limits on the number of MARS requests per user and a limit on the number of tapes used per job?

U. Modigliani replied there is a facility to limit the number of MARS jobs run by any one user and this feature is used occasionally, when one particular user has submitted an excessive amount of work, so that other users' jobs are excluded. The more limits that are imposed on a system, the more likelihood there is of the resources being underutilised, unless monitoring is constant and intervention intensive. He advised users to report to the Call Desk any situations where one individual's work was excluding others'.

Mars documentation difficult to grasp

It is possible to write very inefficient mars requests

U. Modigliani suggested that these two comments were linked. He noted that the documentation advises users to retrieve all parameters for any particular date at one time. He observed that complaints about the MARS documentation were very rare and asked that the user concerned to describe precisely the areas (s)he found unclear. He added that efforts were currently being made to rationalise the location of existing documentation to improve its accessibility. The possibility of rejecting inefficient jobs before they begin running has not been investigated. M. Fuentes noted that the data are organised differently on the different MARS servers, so that a job that runs efficiently when retrieving data from the research database may not run efficiently retrieving data from the operational database.

D. Lucas added that users could monitor the progress or lack of progress of their jobs via the webMARS pages on the ECMWF website.

Difficult to understand which data are available

R. Rudsar clarified that the comment concerned data available via webMARS. U. Modigliani agreed that it was not straightforward but explained that, as new datasets were constantly being added and some datasets were for restricted access, it was difficult to keep documentation up-to-date. The comment was noted and User Support would investigate whether any improvements could be made.

Expver: what is it?

U. Modigliani explained that expver 1 was the main operational stream; it was used implicitly by the data finder and would be what most users would need. It would not be straightforward to provide links to explanatory pages from the long list of expvers. For most purposes, users could remain unaware of the expver number. M. Fuentes proposed that expver 1 could be highlighted and given prominence on the list, as users' first choice by default.

Instability of data finder

U. Modigliani explained that the data finder response times were sometimes unacceptably long. This was due to system overload and was resolved at the beginning of the year. He asked users to report any future instances of poor response times to ECMWF.

ECaccess

- Resilience of job submission, in particular time-critical jobs - will be covered by a presentation later
- High availability features - will be covered by a presentation later. U. Modigliani added that ECMWF tries to maintain a constant process of developments to improve resilience etc.
- Jobs are not sent quickly to ecgate: INIT state for a long time

This refers to a specific situation which arose in a period of heavy overload on ecgate and will also be covered later.

- Display of ECaccess certificate validity for the various commands

U. Modigliani asked for confirmation that what was sought was a single command which would show the expiry dates for the certificates for all the individual services. He stated that this should be a relatively straightforward enhancement to provide.

R. Rudsar asked where the output from time critical jobs could be found, now that they were run under ECaccess. L. Gougeon replied that users can either specify where their job output is to go or the ecjget command can be run as often as required, to retrieve the job output, as the job progresses.

Miscellaneous

- Recursive option in ECFS commands (erm)

U. Modigliani replied that it was still planned to enhance ECFS commands and provide more recursive options, e.g. recursive els and umask have already been implemented; erm will be next, owing to demand from the HIRLAM community, along with echmod and echgrp. I. Weger added that the ecfs client code was being redesigned and would eventually be rewritten to remove reliance on shell functions and make it more modular and thus easier to maintain.

- Users not familiar with Unix
  - Difficulties with strong authentication
  - Difficulties with access to mars
  - Improve web based access to MARS data

U.Modigliani asked whether Spain's requirements could be met by the development of a more powerful, even more user-friendly WebMARS interface. He noted that a data server had been developed for the TGGE project to allow access to various datasets produced by the UK, JMA etc..

- Improve documentation especially for non Unix users

U. Modigliani commented that it had been recognised for some years that new users on the ECMWF training courses were unlikely to have much familiarity with UNIX; whilst ECMWF is unable to undertake basic UNIX training, the structure of the courses tries to take this lack of background knowledge into account. T. Lorenzen suggested that ECMWF could compile a list of basic UNIX manuals for beginners to read in preparation for the courses.

GRIB API: support of fields in rotated lat/long

U.Modigliani replied that basic rotation is supported, in terms of data decoding etc. There is not yet an iterator to navigate the data, so nothing more sophisticated is possible yet. It is being developed.

Provide a common database for the FLEXPART users

The representatives from Spain, Belgium, France, Switzerland, CTBTO and JRC indicated that they would be interested in this facility. U. Modigliani stated that possibilities would be investigated; in particular, whether ECMWF would carry out the development or would aid Member States to organise it themselves.

Possibility to tailor the plots: more fields, different contours, etc

This is being considered as part of the "Service on Demand" project which has developed from the "Plots on Demand" project, though there is no timescale for developments yet.

Earlier availability of 00 UTC model results

A questionnaire on States' requirements for product delivery times was issued recently. Responses are currently being analysed and proposals will be made to the Technical Advisory Committee in October.

Availability of point data for model levels

General availability would be very expensive. Currently point data are available for six pressure levels and some time series (vertical profiles on a point). If users made limited requests for the specific products required, then the cost implications could be assessed.

SMS

I.Weger asked which States were using SMS: Germany, Greece, France, Norway, Finland and Hungary; Belgium, Denmark, Iceland and Ireland are using it or planning to do so.

V. Gislason asked what States were using, if not SMS. M. Andersson replied that SMHI used ControlM because their system used OpenVMS and SMS did not support VMS. (They are now leaving VMS.) P. Halton stated that Ireland used in-house scripts, developed over time. There is currently some discussion about whether it would be worth moving to SMS. R. Barnes noted that the UK had its own suite control system, developed in-house. He believed that SMS was found to be satisfactory for major projects run on the ECMWF computing system.

R. Rudsar noted that Norway had considerable experience in using SMS and had made some developments which could be of use to other services. Would anyone be interested in an SMS Mailing List? U. Modigliani suggested that representatives might like to discuss SMS and compare experiences at their next meeting. ECMWF had not given an update for some time.

V. Gislason would like to install SMS at the Icelandic met. service and would like more information. He was informed that there was documentation on the ECMWF web pages and additional information could be obtained from Iceland's User Support Contact Point. He also asked for some information on GPFS on Linux. He received some initial information from ECMWF analysts and was recommended to keep in contact with them, as they themselves gained experience with GPFS.

E. Krenzien was interested to learn more about the recent extensions to ECFS and was directed to the relevant analysts.

Denmark asked for a command line version of a facility to display MARS queues. *U. Modigliani replied that its feasibility would be investigated.*

U. Modigliani noted that he had also been asked when ECMWF would begin to produce fields in GRIB 2. He noted that once the resolution was increased, in particular in the vertical beyond 127-128 levels, then it might be beyond the range of GRIB 1 and GRIB 2 would have to be introduced; however, this is at least 3-4 years away.

I. Weger noted that ECMWF had undertaken to investigate the feasibility of a common database for Flexpart users.

P. Halton asked that a summary of the information presented at the Security Representatives meeting on IPv.6 be made available to Computing Representatives. *I. Weger agreed that this could be done.*

P. Hitij suggested that the material in the printouts provided to all the participants could be made available via a temporary internal website, so that it could be accessed via notebooks and laptops during the meeting, thus reducing printing requirements.

NEXT MEETING

I. Weger asked representatives whether they had found the new format -of one full day and two half days- to have been successful. There was general agreement that the additional time had been very useful. They also agreed that it was practical to have the meeting back-to-back with the Security Representatives' meeting and that the next meeting should be in twelve months' time. I. Weger stated that the date, in spring 2008, would be finalised in due course.

PART II

Member States' and Co-operating States' Presentations

AUSTRIA AUSTRIA

Dr Gerhard Hermann – Zentralanstalt für Meteorologie und Geodynamik, Austria

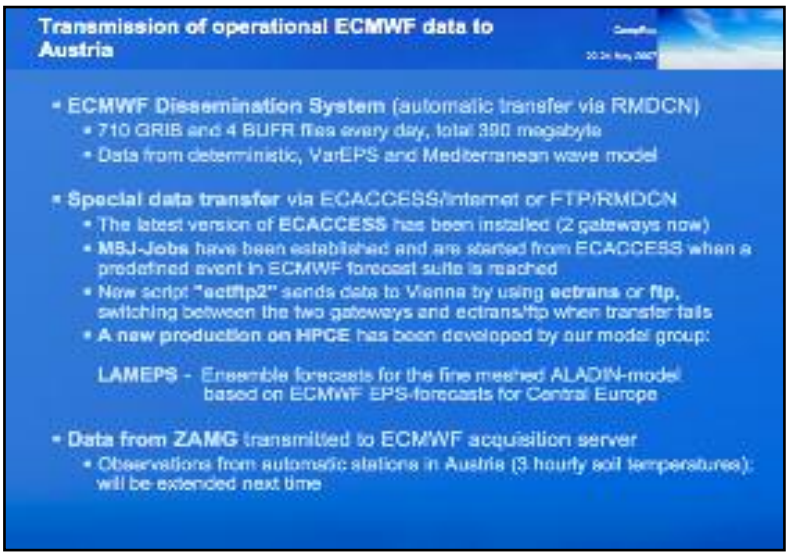

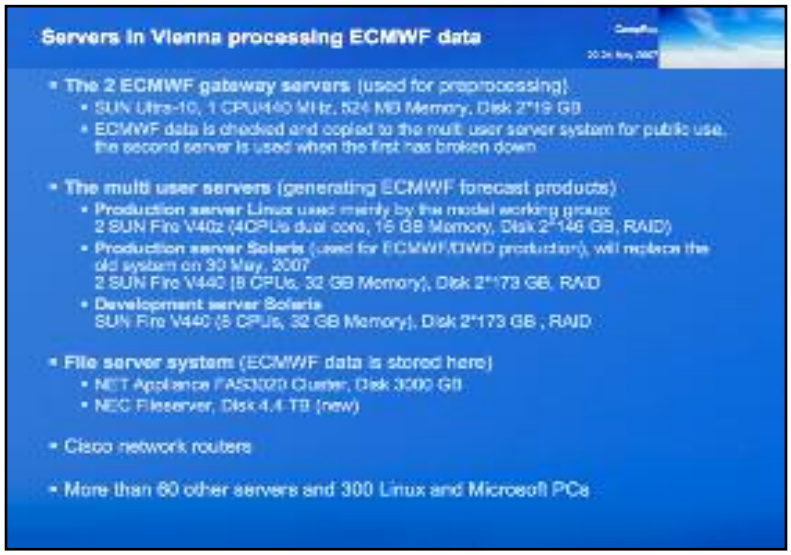

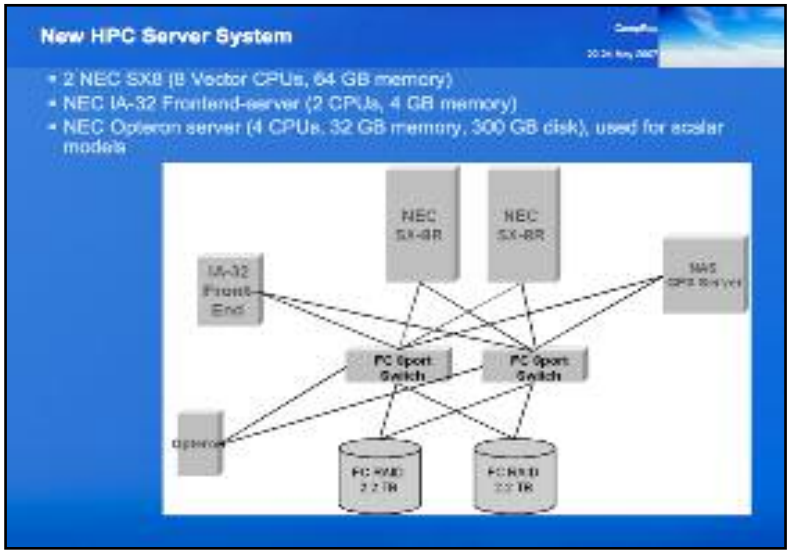

AUSTRIA

AUSTRIA

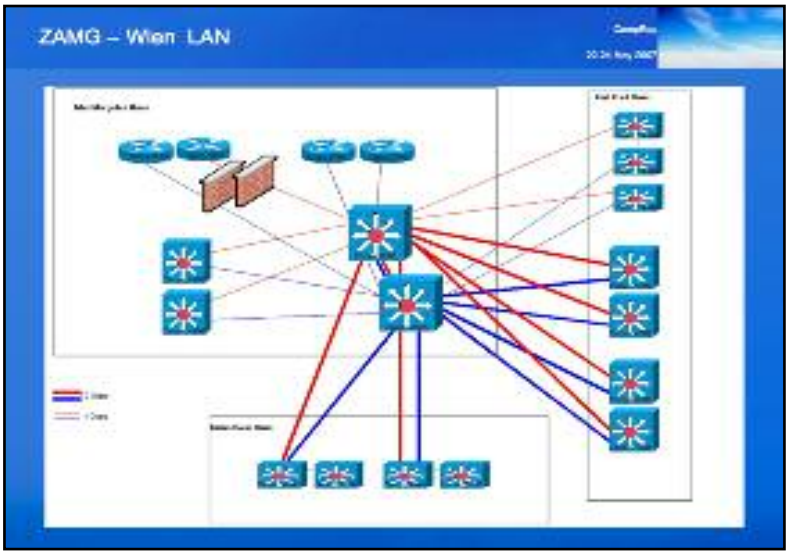

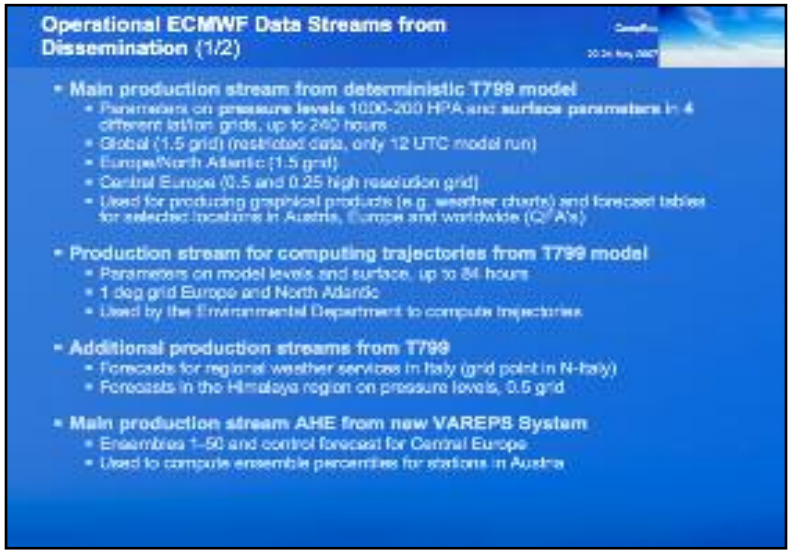

AUSTRIA

AUSTRIA

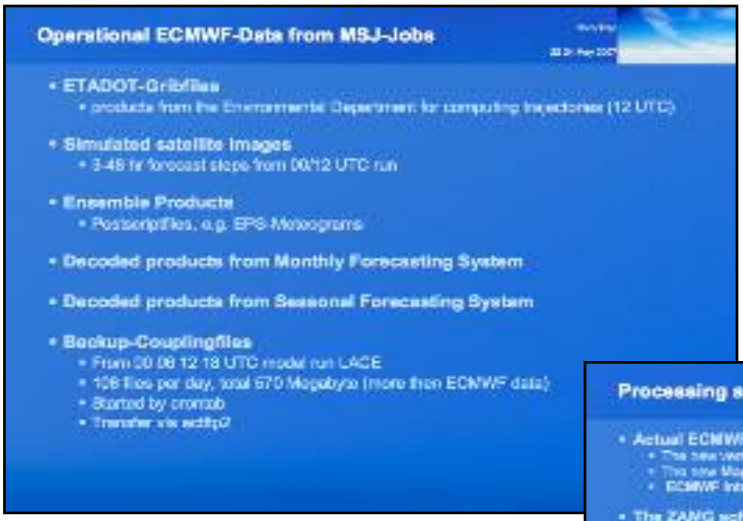

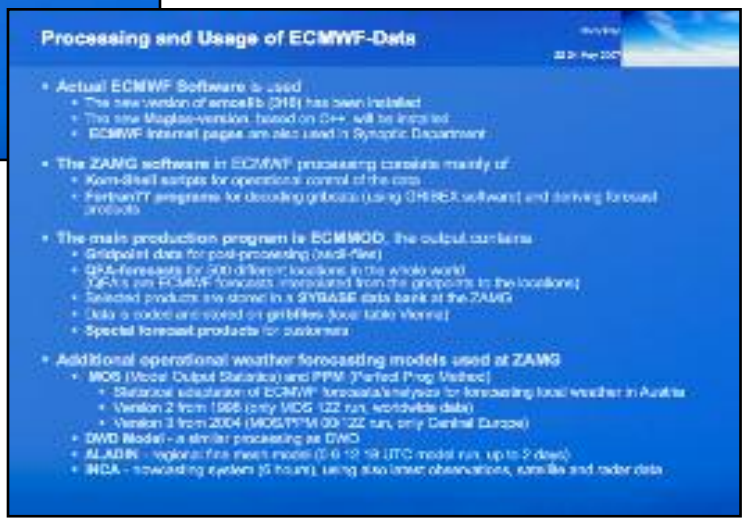

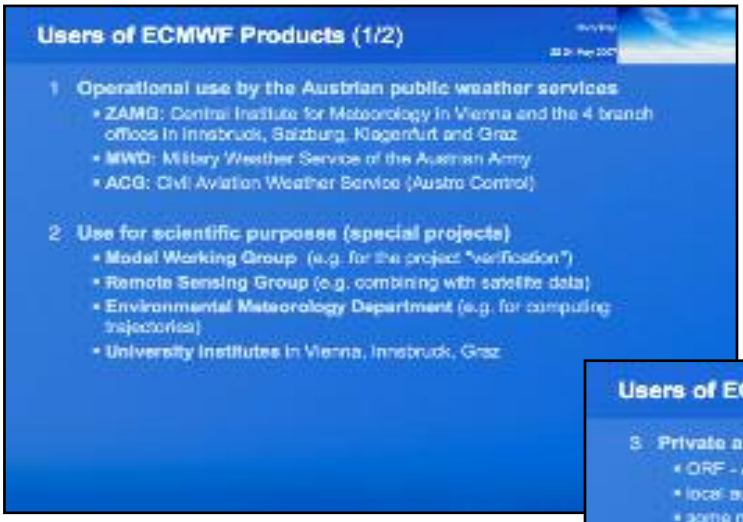

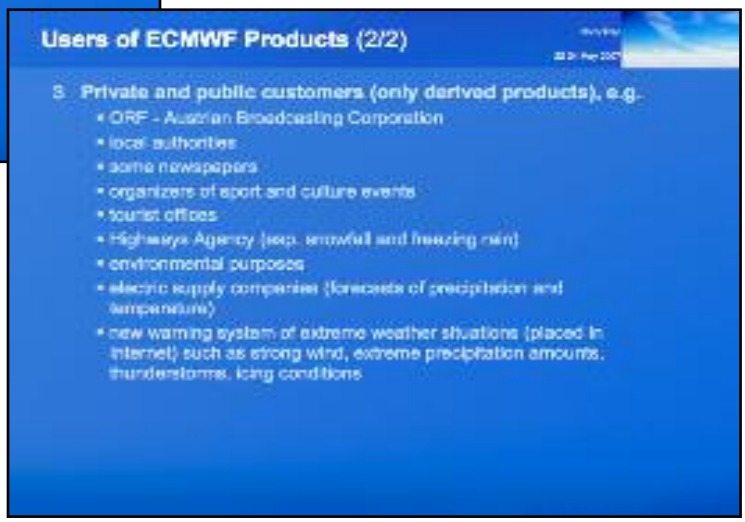

BELGIUM BELGIUM

R. Swennen – Royal Meteorological Institute, Brussels

Computer infrastructure evolution at the RMI of Belgium

Three institutes located on the site of the Plateau d'Uccle (the Belgian Institute for Space Aeronomy (BISA), the Royal Observatory of Belgium (ROB) and the Royal Meteorological Institute RMI) share some common servers: a file server and intensive computing server. Each institute also has a local network with specific infrastructure.

Last year we upgraded our operational Aladin runs to four runs per day, and we changed our central SAN infrastructure from an HP-VA7400 to an IBMDS4800. For the moment we are planning our migration from our central file server HP to the new one a Netapp FAS3050 cluster.

Specific RMI infrastructure:

More and more servers are migrating from HP-UX to Linux. Not all because I believe Linux is not as stable as HP-UX. The main production server, telecom server and the database server are running on HP-UX servers. Additional image production is done on a cluster (Heartbeat) of 2 Dell servers running SLES9. The RMI web and FTP server are running on DELL Power Edge RHEL4. In the near future we will migrate from our HP-UX login server to a RHAS4 DELL login server. Last year a Microsoft Active Directory domain was implemented and there may be a possibility to integrate it with Exchange in the near future.

Communication links:

In the near future our primary Belnet connection will be upgrade to 1 Gbps and our backup connection to Belnet

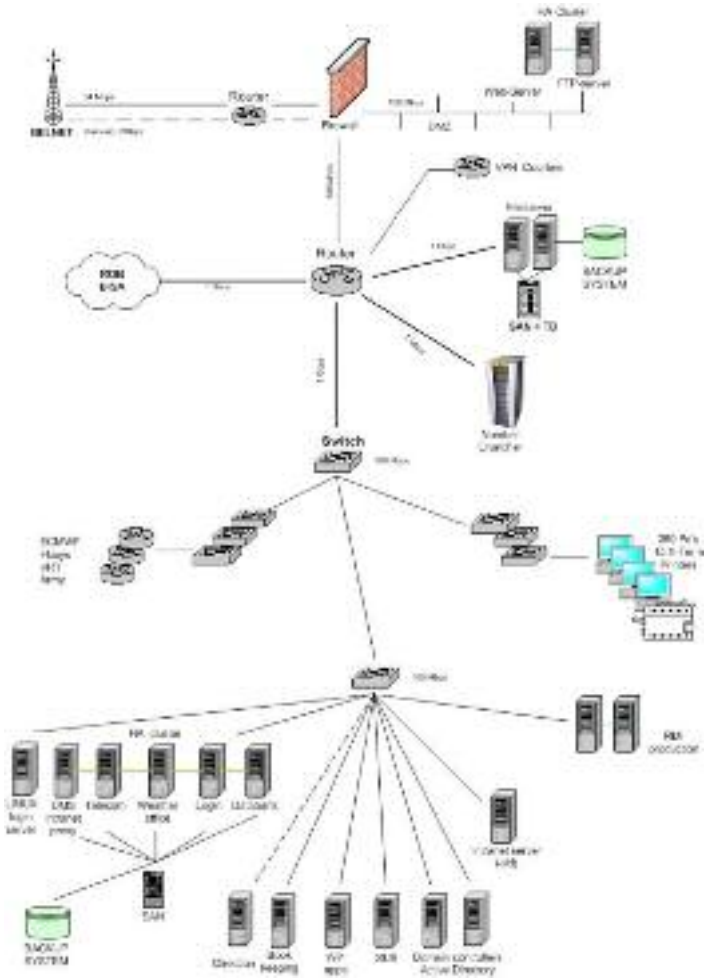

will be established by a 100 Mbps connection. Currently we are migrating the RMDCN network. Our connection infrastructure to the branch offices and customers hasn't changed.

The three institute shared infrastructure:

We bought a Netapp FAS3050 (6TB Fibre channel and 10TB S-ATA) NAS server and are planning the migration from the old file server to this one.

Currently we have beside our parallel compute server 4 "KAOS" machines. They can be used for long mono CPU jobs (Matlab, IDL, ...).

The hardware computer infrastructure of the RMI:

The servers are at least for this year still mainly based on HP servers with progressive introduction of Linux servers:

- HP 9000/L1000 (2 processors, 1 GB memory, 34 GB internal disks): oracle database server.
- SAN infrastructure consisting of an IBM DS4800.
- HP 9000/L2000 (4 processors, 2 GB memory, 72 GB internal disks): login and application server. In a replacement phase.
- Power Edge 6950 4 Dual Core Opteron 16GB memory machine, will replace the HP-UX login and application server.
- HP 9000/L1500, 2 processor, 512MB memory and 34 GB internal disks: telecommunications server. All our meteorological products (RMDCN) are sent and received on this server.
- HP rp3440 2 processors, 1 GB memory, 34 GB internal disks): forecast office server.
- HP rp3440 (1 processor): used as web applications server (Tomcat), proxy server (Squid). It runs also several softwares such as a DMS (Apache Jakarta Slide), workflow (OSWorkflow), access to Oracle data base (Hibernate), server monitoring (Bigbrother), problems logging (Elog, free ware using a ticket system).

These HP-UX servers are tied in a high availability cluster.

CROATIA CROATIA

Vladimir Malović – Croatian Meteorological and Hydrological Service

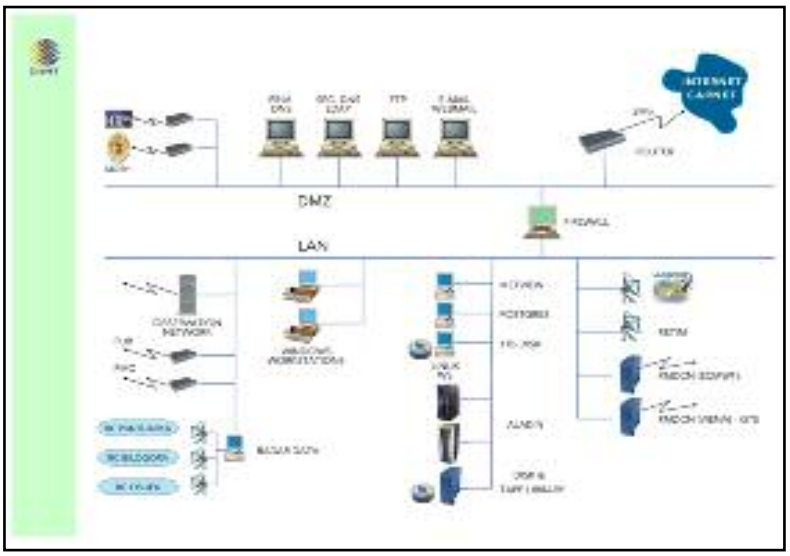

Computing resources

- 16 processors SGI Altix 3000, 32 GB RAM (irix)
  - LAM Aladin, Regional Climate Model (RegCM)
- 24 processors SGI 3400, 48 GB RAM
  - SuSE linux
  - LAM Aladin, RegCM
- SGI Origin 200
  - GTS, e-mail, radio sonding, climatological db, radio sonding db
- SGI Origin 350
  - Qualstar tape library (5 TB online)
  - 1,5 TB disks space

Computing resources (2)

- LINUX servers
  - MetView (version 3.9)
  - PostgreSQL database (climatological data)
  - 1.5 TB disks space
  - TnVis (visualisation for TV)
  - EUMETSAT
  - RETIM (Meteo France satellite dissemination, GTS)
- MS Windows computers
  - DWD data
  - Observation network and automatic station data

Future plans

- VPN network (in progress)

-Remote offices at coast
-Connect Radar centres
-Replace ISDN lines

- Var EPS – products to D+15 (in progress)
- Upgrade tape library to 14 Tb, disk space up to 13 Tb

CZECH REPUBLIC CZECH REPUBLIC

Karel Ostatnicky – Czech Hydrometeorological Institute

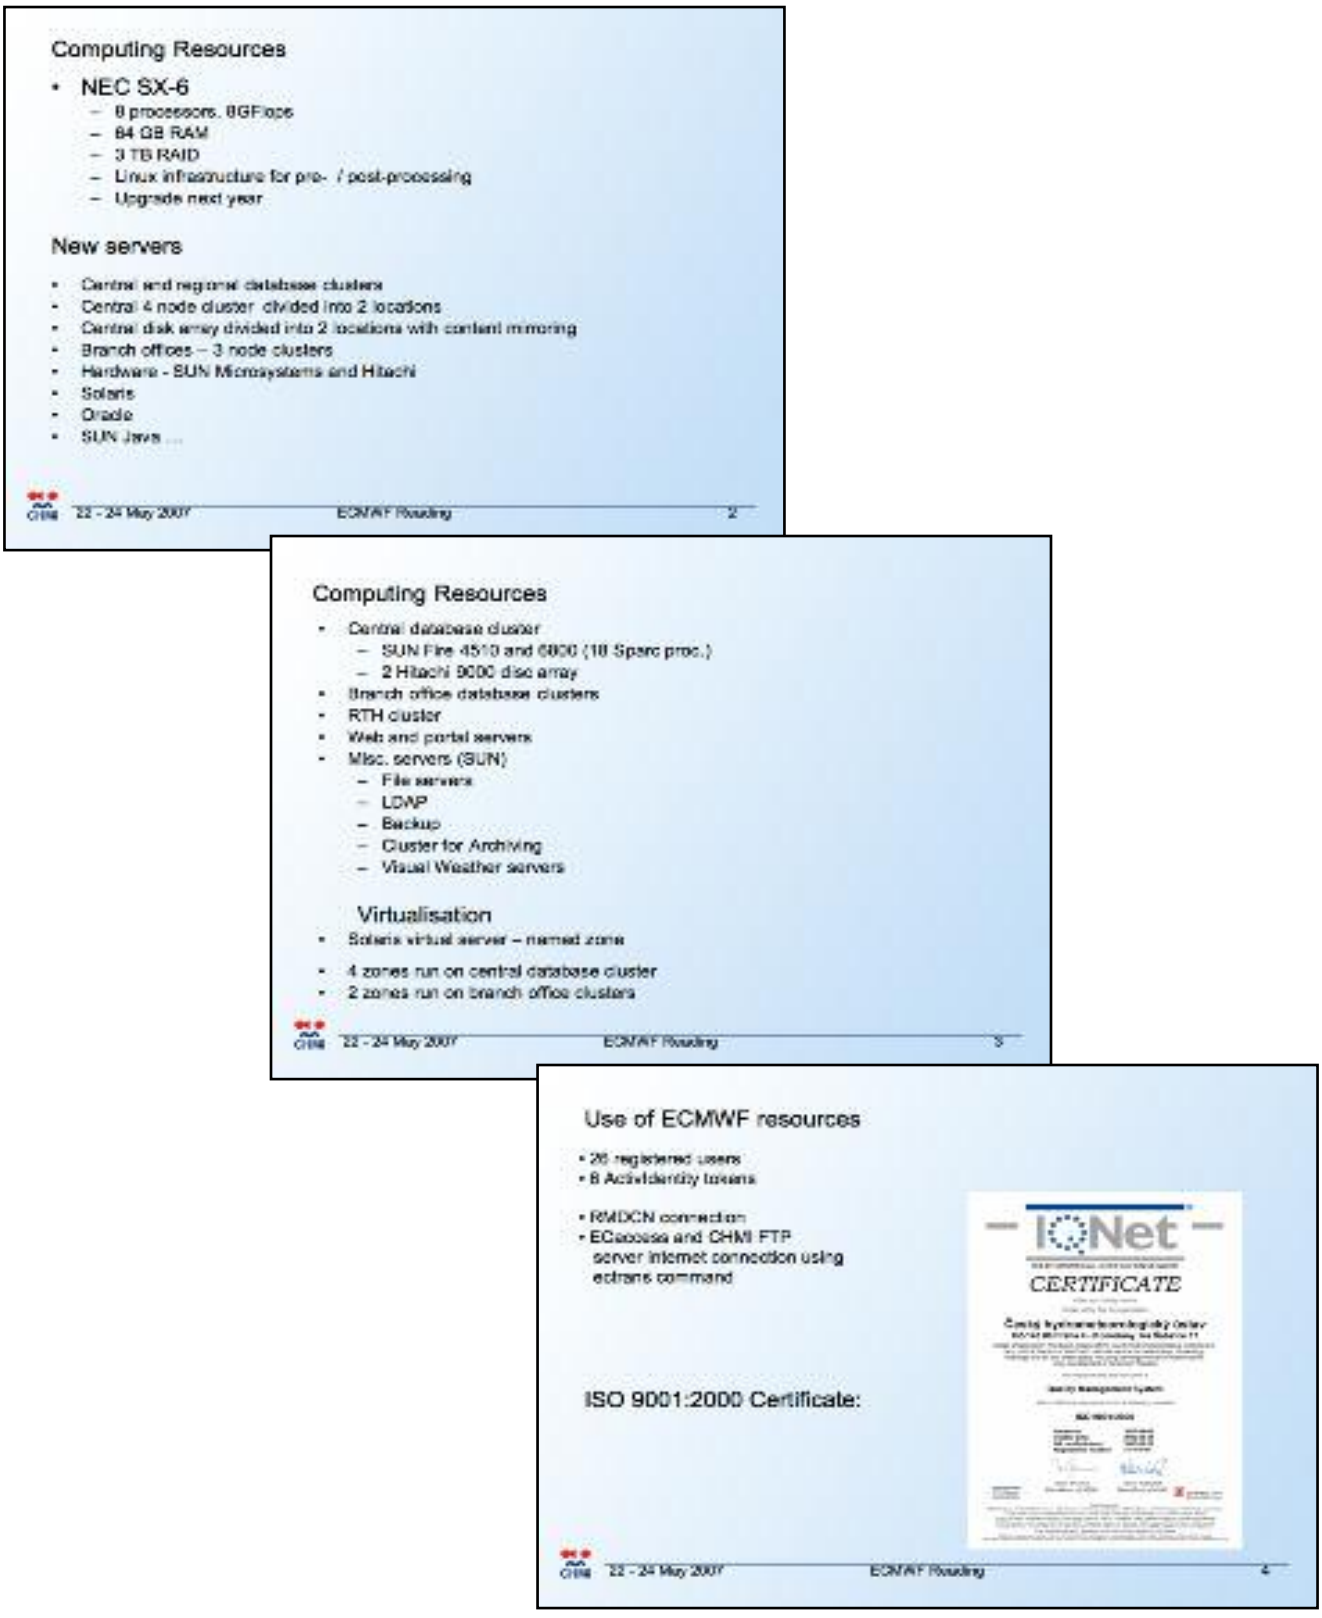

T. Lorenzen asked how they implemented and controlled the content mirroring of the central disk array. *K. Ostatnicky replied that they used software supplied by Hitachi.*

DENMARK

DENMARK

Thomas Lorenzen – Danish Meteorological Institute

DMI local computing facilities
The current HPCF in brief

- NEC SX6 64M8 (8 vector nodes with 8 CPU's each)
  - 8~ (8~ 8) Gflops
  - 8~ 32 Gbyte RAM
- NEC TX7 16M2 (2 scalar nodes with 8 CPU's each)
  - 2 ~ (8 ~ 1.30) GHz Intel ItaniumII
  - 2 ~ 16 Gbyte RAM
- SX6 is used for running the Hirlam weather forecast model four times a day.
- TX7 does pre-processing and post-processing of Hirlam data plus running of demanding multi processor scalar jobs for air quality and ocean modelling.
- Other operational workloads, among which also GTS receiving and archiving from RMDCN, are done on quite a big park of Linux and Solaris based machines.
- Network communication is based on 1000Mbps TCP/IP in the backbone and to selected hosts with 100Mbps TCP/IP for the rest.

DMI local computing facilities
Current Hirlam forecast areas and resolutions

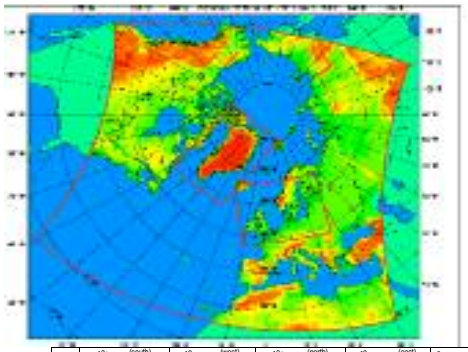

	T15	S05	Q05
^{n}lon	610	496	550
^{n}lat	568	372	378
^{n}vert	40	40	40
Resol.horisl	0.15°	0.05°	0.05°
Δt	450s	150s	150s
Forecast length	60h	54h	36h
Host model	ECMWF	T15	T15
Boundary upd	3	1	1
^{n}Output freq	1	1	1

	$l^{a}t_{min}$ (south)	$l^{o}n_{min}$ (west)	$l^{a}t_{max}$ (north)	$l^{o}n_{max}$ (east)	$l^{a}t_{south\ pole\ of\ rotation}$	$l^{o}n_{south\ pole\ of\ rotation}$
T15	-37.527°	-64.325°	47.523°	27.025°	0.000°	80.000°
S05	-1.027°	-13.674°	17.523°	11.075°	-40.000°	10.000°
Q05	-8.000°	-32.200°	10.850°	-4.750°	0.000°	50.000°

DMI

DMI local computing facilities
Procurement for a new HPCF

- The procurement for a tenfold compute capacity increase runs as a restricted procedure.
  - Notification of tender plans in EU journal with possibility for everybody to express interest.
  - After the interest notification deadline, DMI selects 5 vendors to tender.
  - The 5 vendors submit their tenders before the tender deadline (Wednesday noon this week).
  - DMI evaluates the tenders and selects a successful candidate for contract negotiation and signature.
- Milestones in 2007 and 2008.
  - 2007-03-30 : Five companies were selected and forwarded the tender material.
  - 2007-05-23 : Deadline for the five companies to deliver their tender.
  - Summer 2007 : Tender evaluation.
  - Autumn 2007 : Vendor selection with contract negotiation and signature.
  - Either : Installation in one phase, if power consumption not too large for current HPCF to co-exist.
    All to be completed and accepted 2008-12-31 at the latest.
  - Or : Installation in two phases, if power consumption too large for current HPCF to co-exist.
    - Parallel running 1st phase with the current HPCF during code migration.
    - Installing 2nd phase after removal of the current HPCF.

    All to be completed and accepted 2008-12-31 at the latest.

DMI

DENMARK

DENMARK

DMI local computing facilities
Plans for Hirlam adjustments on the new system

- The big T15 domain (0.15° ×0.15°) will be a T10 (0.10° ×0.10°) or better domain of the same geographic coverage.
- The nested domains S05 and Q05, based on the hydrostatic Hirlam model, will be replaced by the non-hydrostatic Aladin model in better horizontal resolution.
- Transitioning from 3DVar to 4DVar.
- Implementing a Hirlam or Aladin based ensemble forecasting system.

dmi

Use of ECMWF facilities
Current utilization of the GTS RMDCN line capacity

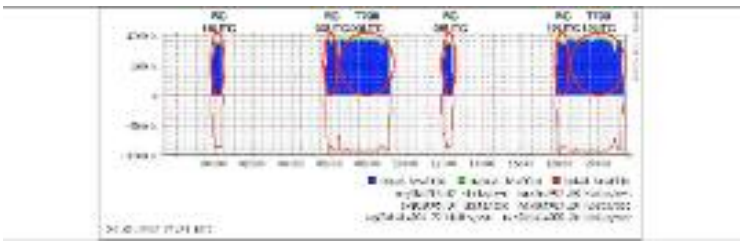

dmi

Use of ECMWF facilities
Operational dataflow from ECMWF to DMI

- DMI receives from ECMWF via RMDCN data amounting to close to 2.5GBytes per day.
- Frames for hirlam account for ¼ of the data volume. Global gauss surface fields for web and other visualization purposes account for ½ of the data volume. The remaining fraction contains a mixture of products for other purposes.
- Around dissemination hours the primary 1Mbit RMDCN line is fully saturated.
- The backup ISDN line only holds one third of the bandwidth of the RMDCN line, so a failure of the primary line will cause operational delays at DMI. DMI may opt for a true mission critical setup, when RMDCN turns to use MPLS.
- For testing and quasi time critical activities, DMI receives from ECMWF via Internet data amounting to about 1.0GBytes per day.
- Data is transferred via Internet via plain ftp as well as via our ecaccess gateway. Using our ecaccess gateway for this purpose has been ongoing for now two year with no problems.

dmi

DENMARK

DENMARK

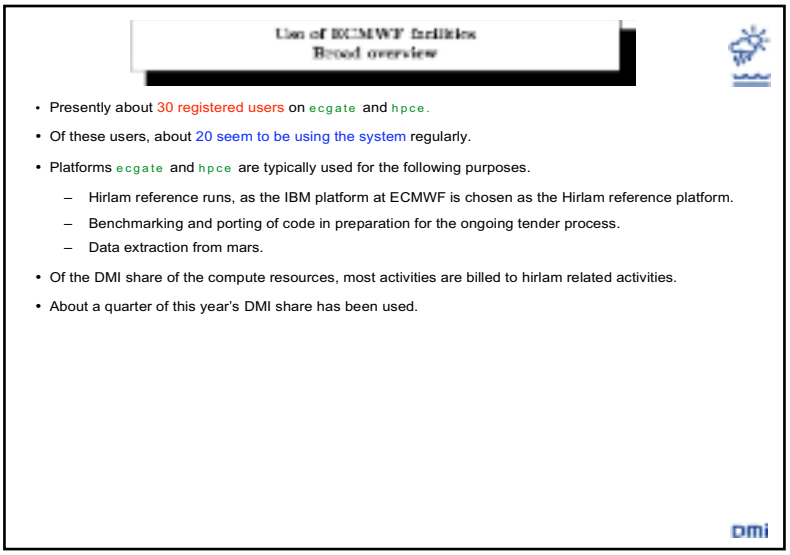

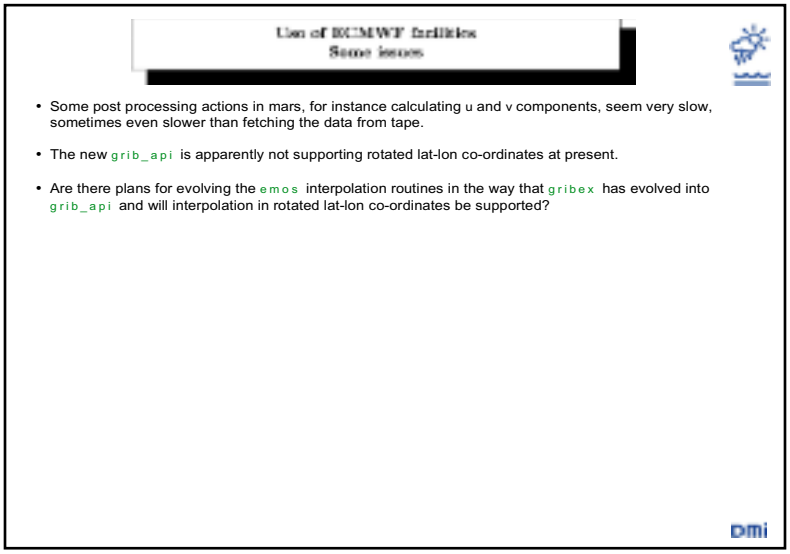

FINLAND FINLAND

Mikko Aalto – Finnish Meteorological Institute

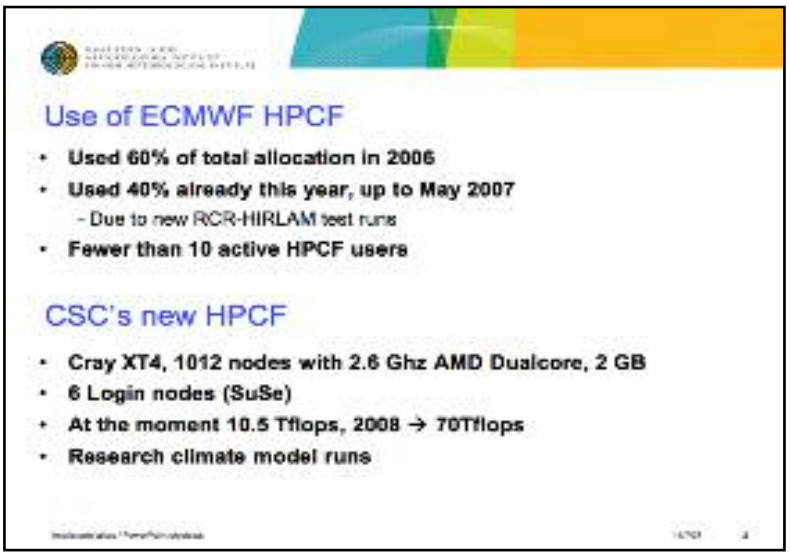

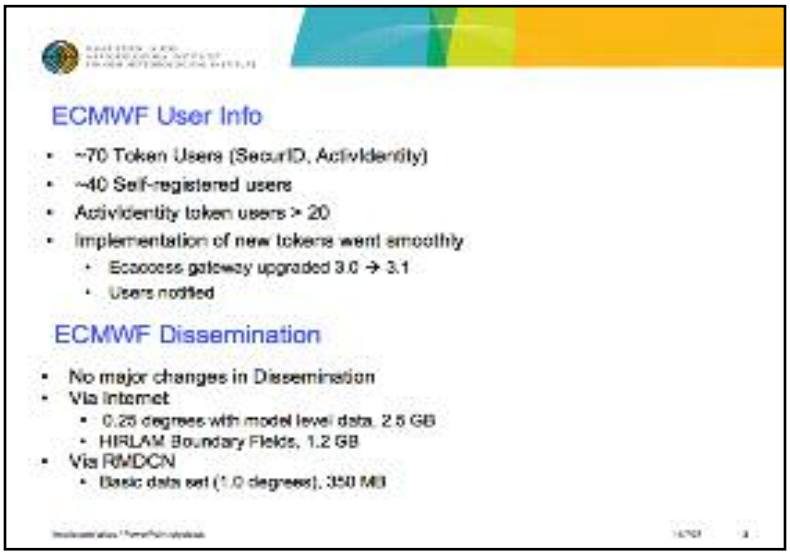

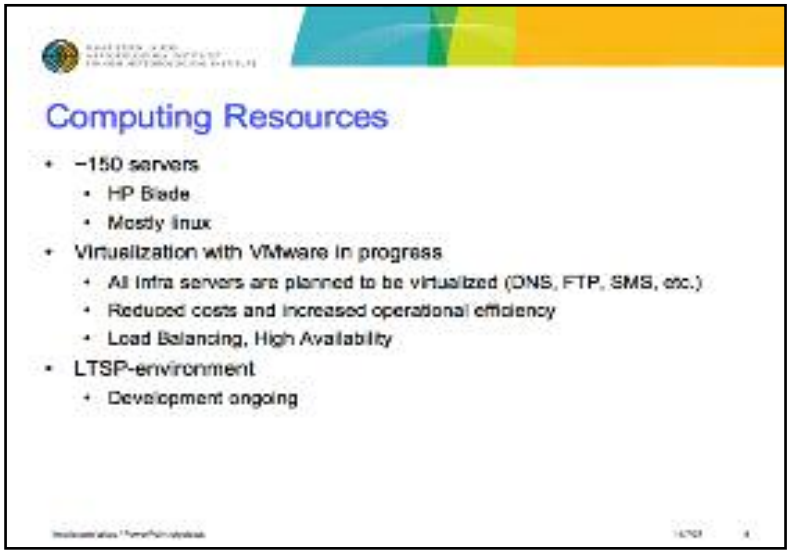

FINLAND

FINLAND

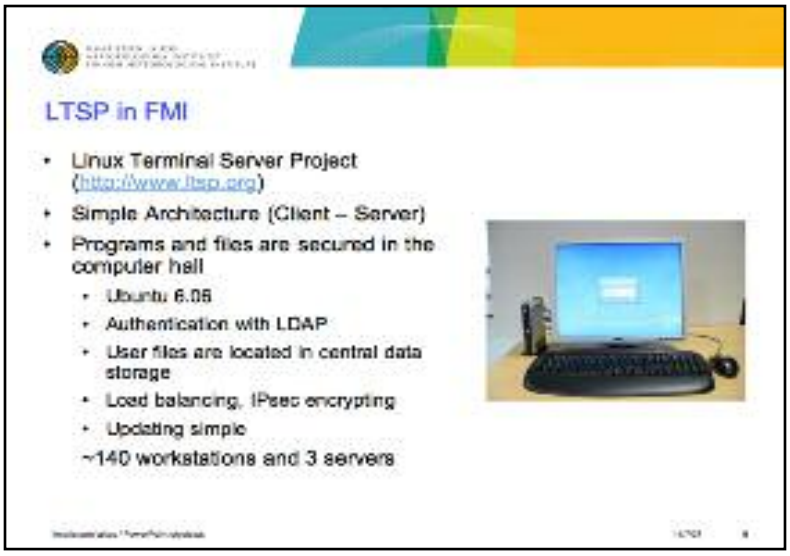

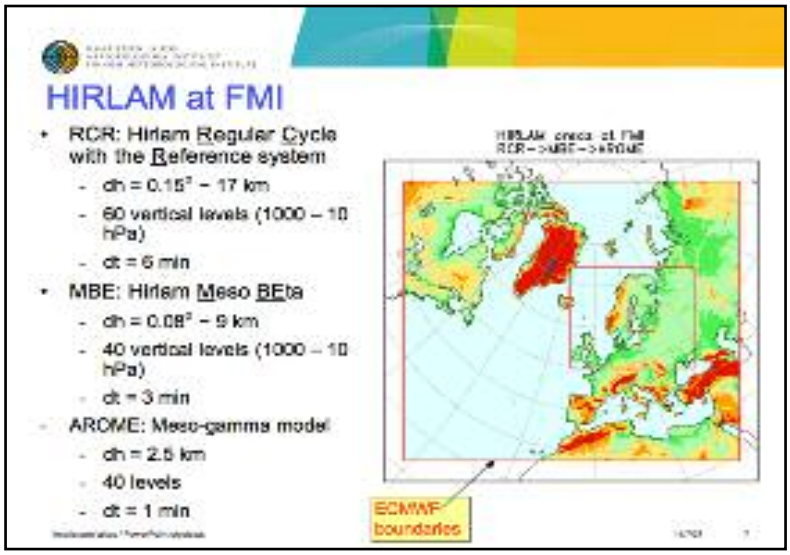

E. Krenzien asked whether they had already tried running Hirlam on the new CSC Cray. *M. Aalto replied that they had not.*

FRANCE FRANCE

Dominique Birman – MétéoFrance

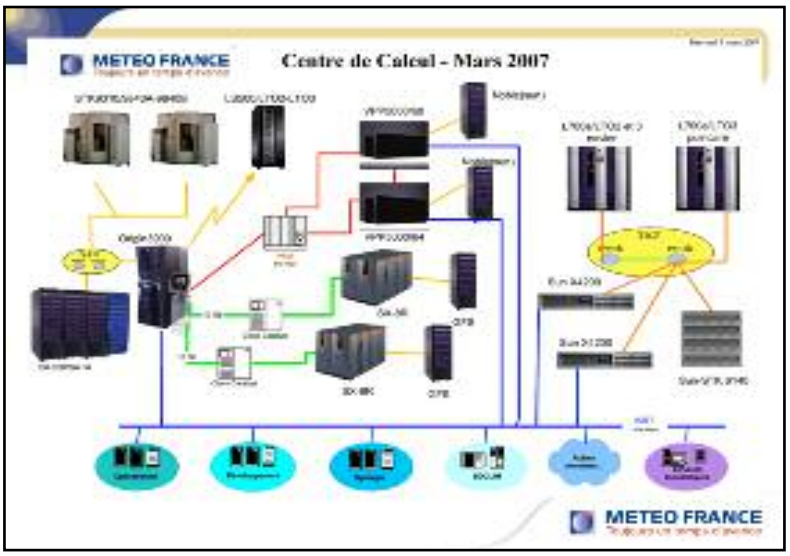

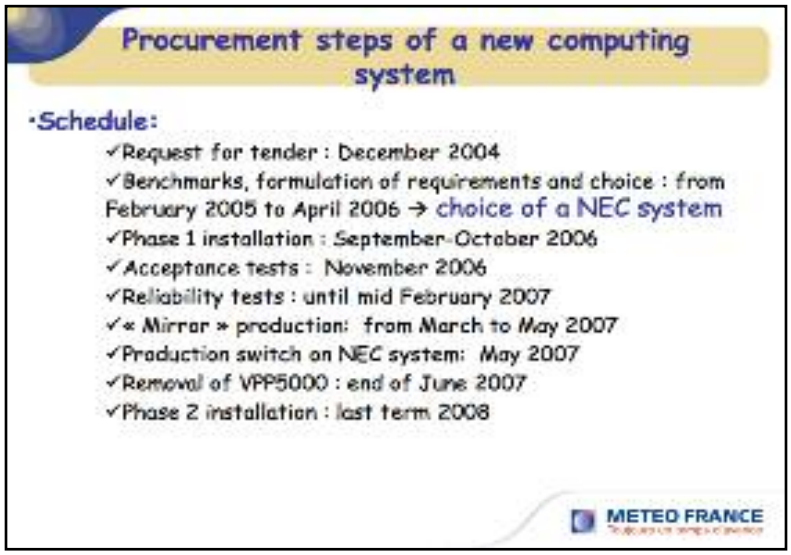

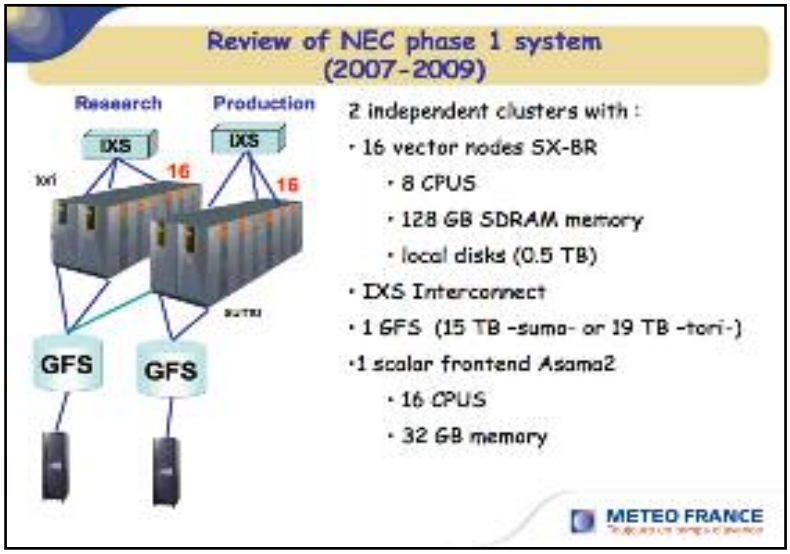

FRANCE

FRANCE

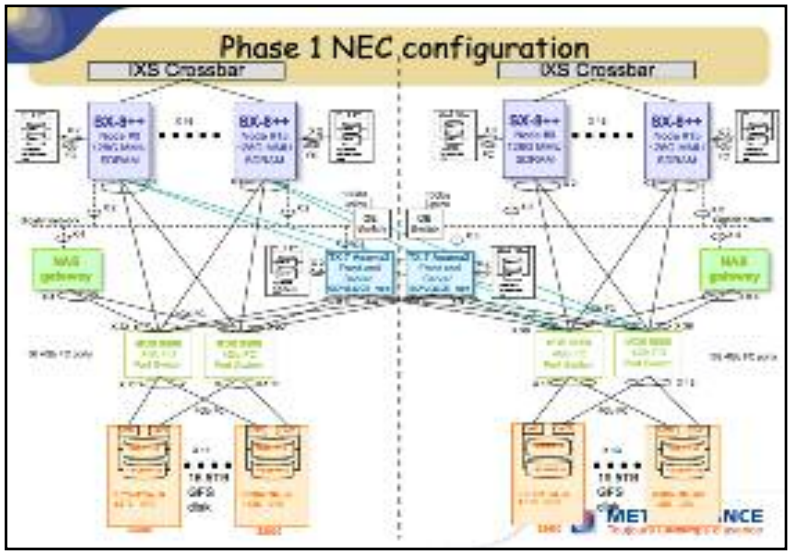

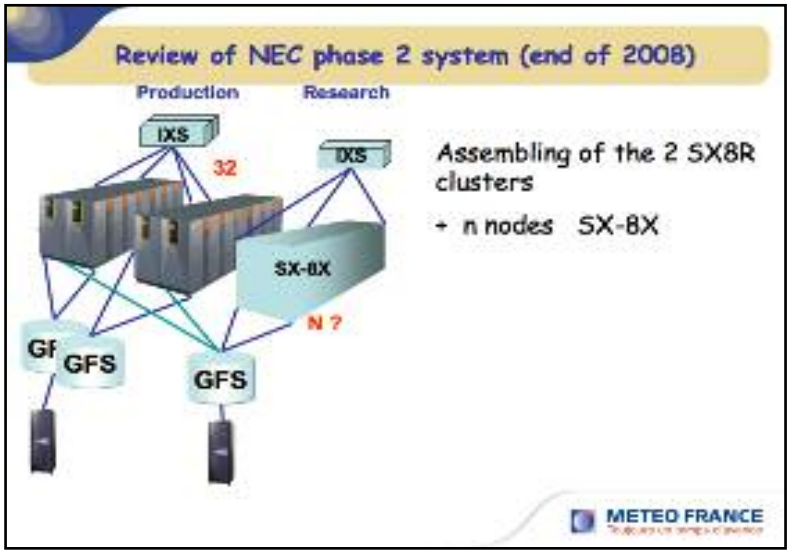

FRANCE

FRANCE

Review of NEC phase 2 system

Phase 2

- Add a new system with new generation SX nodes.
- No removal of phase 1 system (assemble the 2 SX-8R clusters)
- Heterogeneous and non symmetric systems for phase 2.
- Choice of the function given to each cluster (operational/research).
- Suitable front-end systems (powerful enough to cope with the extra load from the new cluster)
- Disk space will be increased by a factor of 4.

METEO FRANCE

Use of ECMWF facilities

- 54 M.F. projects + 4 Special Projects:
  - ✓ 270 users
  - ✓ 141 from Meteo France - 129 from universities or other organisations.
  - ✓ MARS data base success, participation in European projects and use of HPCE resources.
- 13 projects used HPCF resources in 2006 : 8,900 kSBU (81%)
  - ✓ Mercator, Climate, research in NWP, IPSL, LPCE, forecast monitoring, LMD, Benchmark, DIRRE, MESONH
- Special Projects (4) : 67 kSBU (57%)
  - ✓ Coupled ocean-atmosphere models, seasonal forecast (CERFACS), Aladin GMAP co-operation.

METEO FRANCE

HPCF CPU utilisation

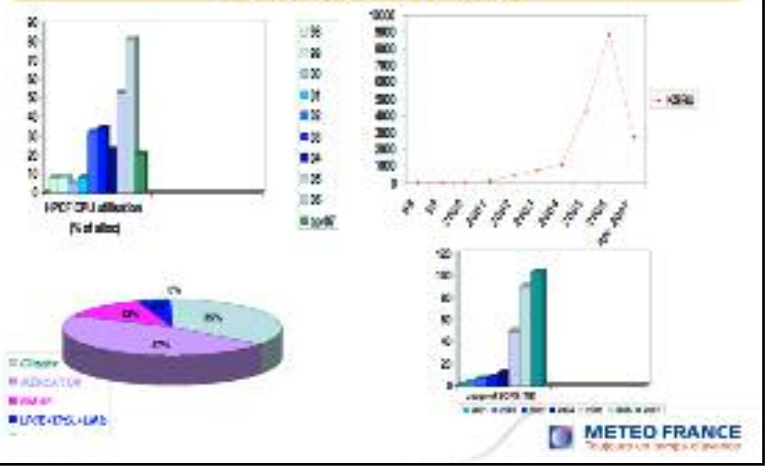

METEO FRANCE

FRANCE

FRANCE

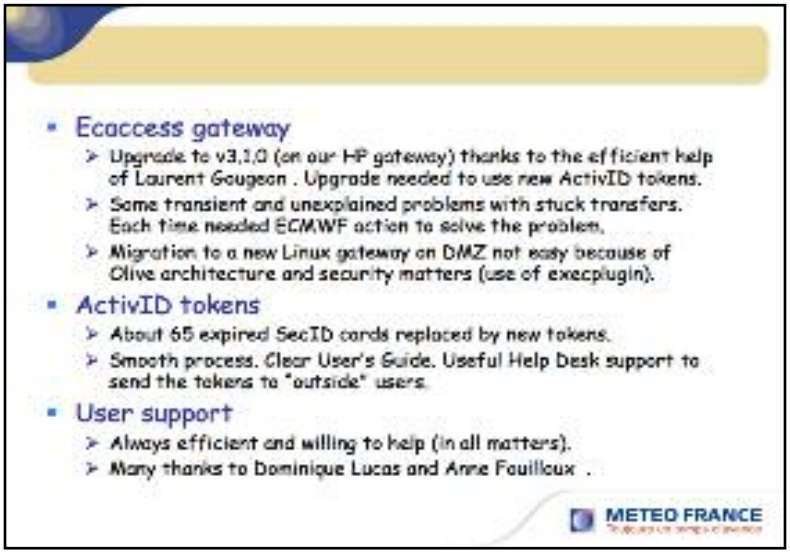

T Lorenzen asked what the front-ends to Météo France's NECs were used for. *D. Birman replied that they were used purely for job submission and interactive access, no operational applications are run on them.*

GERMANY GERMANY

Elisabeth Krenzien – Deutscher Wetterdienst

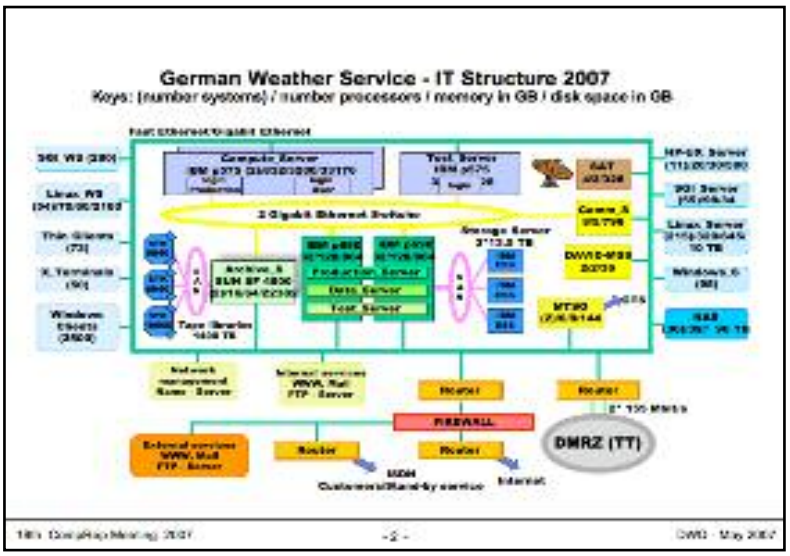

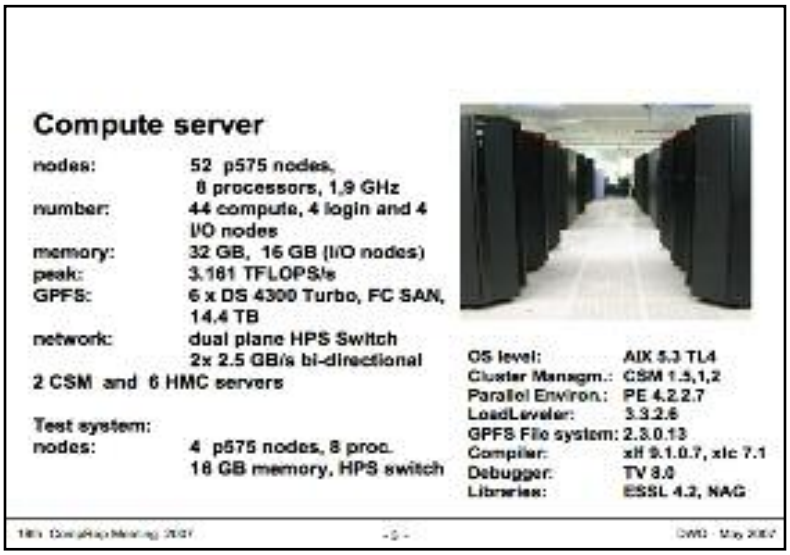

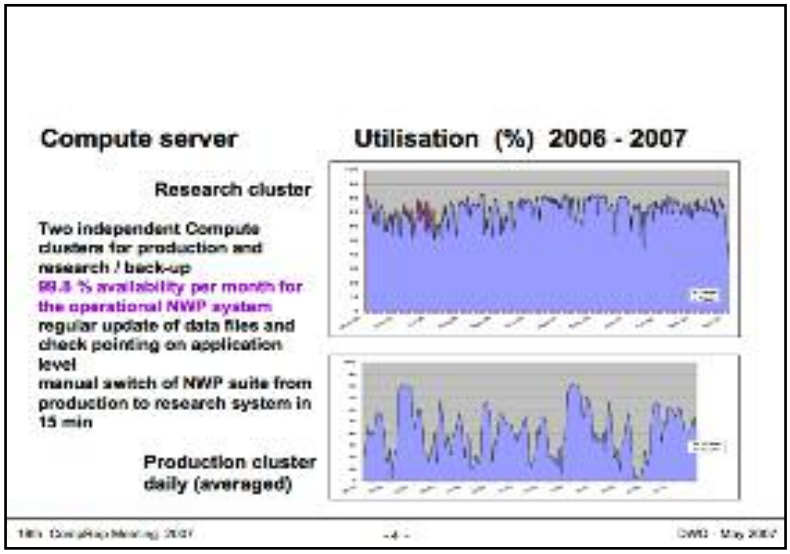

Archive system SAM-FS / QFS

Server:
Failover Cluster: 2 x SunFire 4900 Server
Cluster element: 8 processor US IV, 32 GB memory
Test Cluster: 2 x SunFire 490 Server, 2 Processor 8 GB memory
Storage:
22,7 TB disk cache, Sun StorEdge 6130 system, 140 x 146 GB disks
5 x SAN FC switches for STK 9310 ACS (3 Silos with 17000 slots)
46 FC-tape drives: 16 x 9840 A (20 GB, 10 MB/s)
22 x 9940 B (200 GB, 30 MB/s)
Software: 2 x T10000 (500 GB, 120 MB/s)
Solaris 9 Update 7
Sun-Cluster 3.1
SAM-FS / QFS 4.5.35
ecfs (user interface, ECMWF)
dmscp2 v1.0.7 (data transfer software, M. Stolle, ZIB)

Archive - Tapes

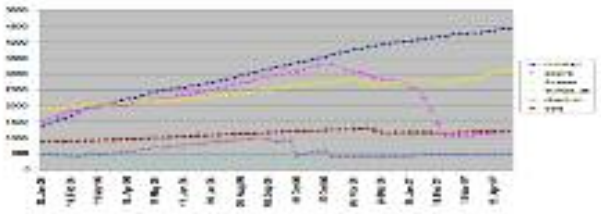

number of files: 12,3 mill.
Total amount: 1.4 PB daily growth rate: ~ 3 TB

Current operational NWP Suite

GME:
Δx = 40 km, 368642 * 40 GP
Δt = 133 sec, Ttot = 7 days

LM-E:
Δx = 7 km, 665 * 657 * 40 GP
Δt = 40 sec, Ttot = 78 hours

LM-K:
Δx = 2.8 km, 421 * 461 * 50 GP
Δt = 30 sec, Ttot = 18 hours

Operational timetable of the DWD forecast models

GERMANY GERMANY

Status and Plans (1)

New building in Offenbach
Cornerstone: 23. March 2006

Moving in: spring 2008
Computer hall: start Dec. 2007
and March 2008

- Move operations to new building in 2008
- Running ensemble predictions based on LM-K in early 2009
- Increase of compute power by factor of 15
- Merge both parts of DMRZ in Offenbach

Status and Plans (2)

Replacement of main DMRZ systems:

- ITT Compute server and Data Handling system:
  16.03.2007 - 15.06.2007
  contract signed: November 2007
  begin of installation: January 2008
  acceptance test: March 2008 (start of migration)
  parallel production: April – June 2008
  decommissioning P4/P5 systems early July 2008
- ITT Tape library (STK 9310 ACS silos):
  June/July - August/September 2007
  contract signed: November 2007
  begin of installation: March 2008
  decommissioning STK 9310 ACS end of 2008

ECMWF services: DMRZ Operations

ECFS enhancements:	AIX 5.3 client
	early-access tests of ECFS_FTP Portal (ECcmd based)
ECaccess v3.1:	Internet: Linux (DMZ), LAN: AIX
	RMDCN gateway: still planned
	Csomars version 1.2 in full operation
SMS v4.4.5:	single instance, solely used for operations
emoslib v310	request for research cluster
Operational Products: (per target date)	via dissemination 657 MB; (transfer rate ~43 kB/s)
	via Csomars 717 MB; (transfer rate 115–1540 kB/s)

GERMANY

GERMANY

Deutscher Wetterdienst

ECMWF resources: Research Modelling

Global Model (GME):	2 reference runs (daily) for DWD production
GME-EPS:	testing assimilation schemes for next model (ICON)
Local Model (LME):	testbed for model improvements
COSMO LEPS:	LM Local area Ensemble Prediction System
EPS studies:	MUSTOK Project 'Storm surges in the Baltic Sea'

	GME	GME-EPS	LME	MUSTOK	COSMO-LEPS
2006	40 km/L40	220 km/L40	7 km/ L40	T_L399/L62	10 km/L40
	20 km/L50	M32 to M64	3TL, 2TL	M51 M10	
ESPC	4200	1000	3300	2000	1900
2007	40 km/L40	220 km/L40	7 km/ L40		10 km/L40
	20 km/L50	M32 to M64	2TL, 2TL		

GME – EPS First results

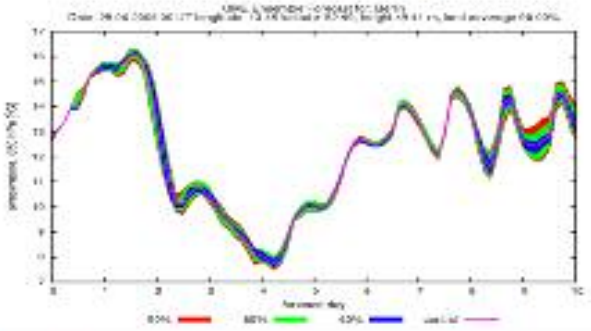

Plans

LMK-EPS	Testbed for research (25 members, 25 kunits per run) implementation has started
CM-SAF	'Pilot' installation of specified components of the software pool on HPCF and Linux cluster
DMRZ-Backup	Implementation of backup service for DWD's operational NWP (scenario: a major outage of the DWD site) begin: 3Q 2007 end : 3Q 2008

GERMANY

GERMANY

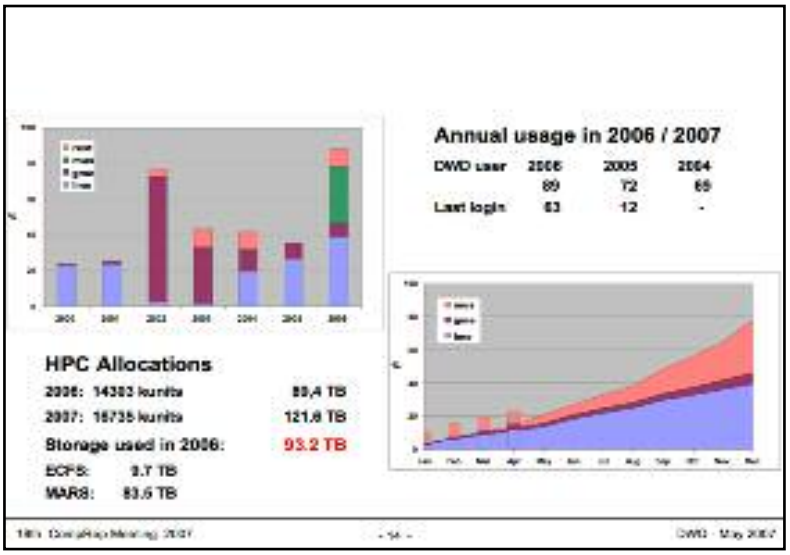

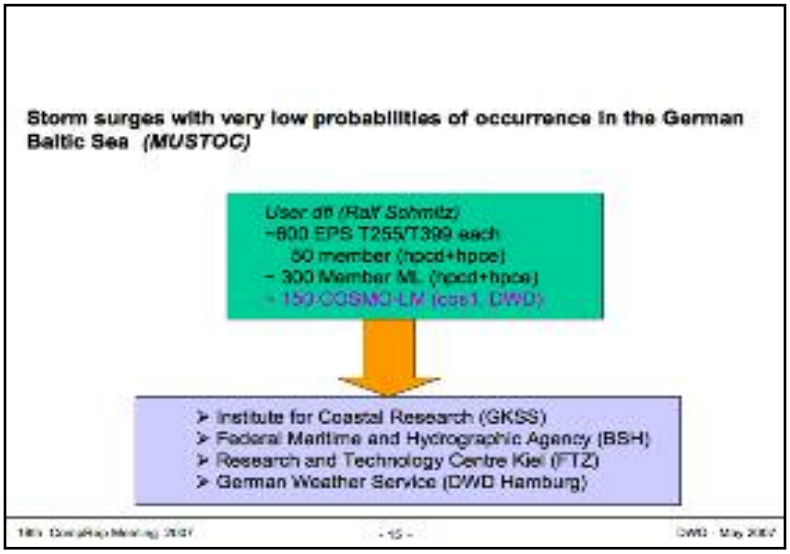

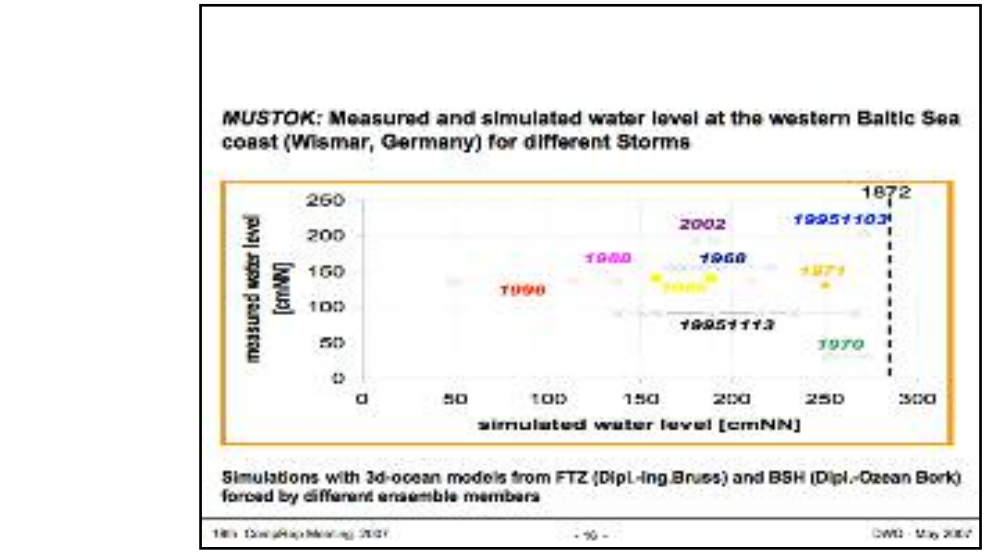

GERMANY

GERMANY

Special Projects (SP)

Currently 28 SP are registered for German universities (12) and different research facilities such as Max-Planck-Institutes (MPI, 5), DLR (4) and other research centres (7) and 3 new applications for 2008 are handed in.

number of users:	2007 (April)	2006	2005
	101	87	81
last login	63	17	-

Main activities:
HPC Resources are used by 12 SP mostly for large scale climate modelling
Retrieval of MARS data
Current forecast data for flight guidance during field campaigns (several DLR and research institutes)
ERA-40 data access

NinJo key features

- Partners: DWD, BGIS, MeteoSwiss, Danish Met. Institute, Met. Service Canada
  licensing possible (2 contracts, more interest)
- operationally introduced at DWD & BGIS
  ~150 dedicated workstations, ~60 servers
- pure Java Client – Server system
  ~ 40 software developers (2 locations)
  ~ 1.5 mio lines of code (Java)
- multi-window technique
  main and 0…3 secondary scenes (main window) meteograms, diagrams (e.g. sounding), separate cross-section visualization, table-view, separate
- highly configurable (via XML)
- more details: http://www.ninjo-workstation.com

- layer – oriented visualization (one layer)
  - satellite data, radar data
  - observational data (incl. point forecasts)
  - model forecast data (~ 10 models available)
  - lightning data, trajectories
  - aviation-specific data, road weather data
  - etc. in total ~ 30 layers
- additional components
  - warnings (automatically generated)
  - interactive graphical editing
  - on-screen analysis – pre-operational
  - batch production of graphical products
  - archiving, playback of weather cases

NinJo Screen shots

Anabel, according to my calculations, this should fit on p.3 of the German bit. If not, or too small to be useful, then just omit it. PP

GREECE GREECE

Antonis Emmanouil – Hellenic National Meteorological Service

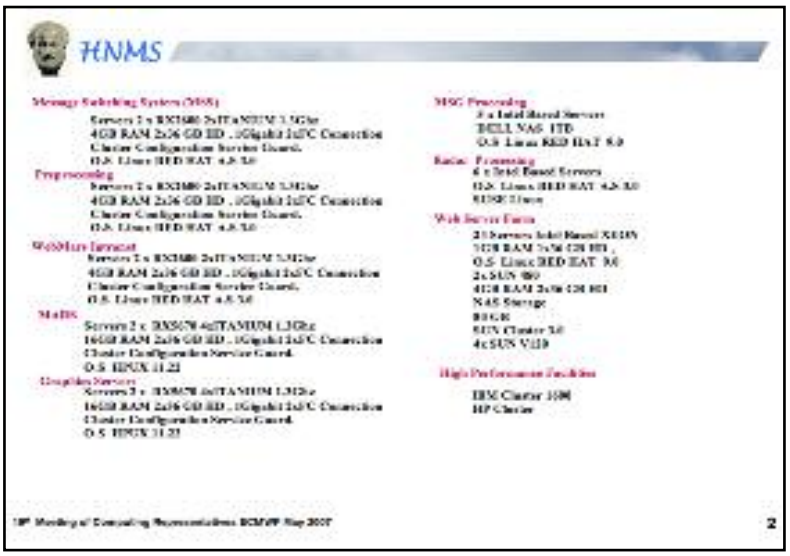

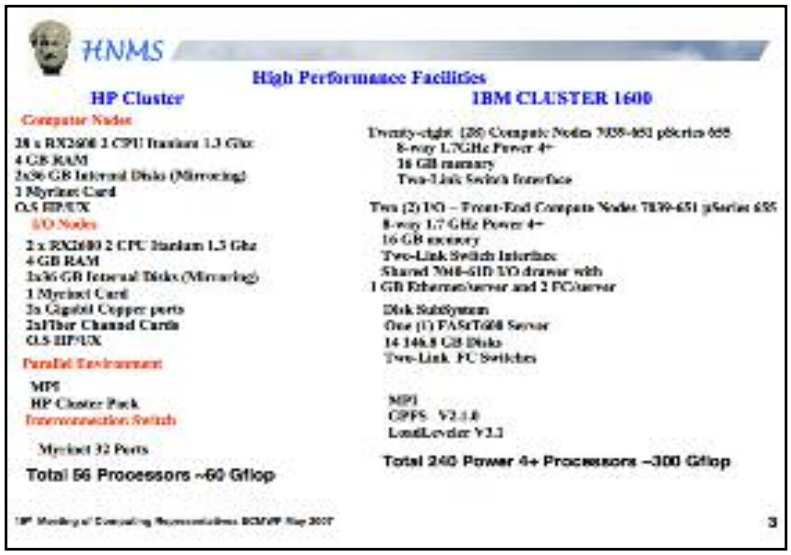

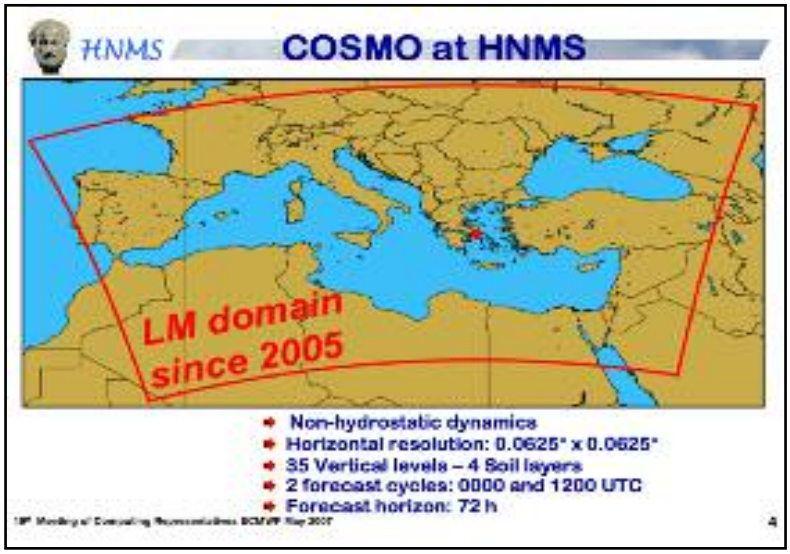

GREECE

GREECE

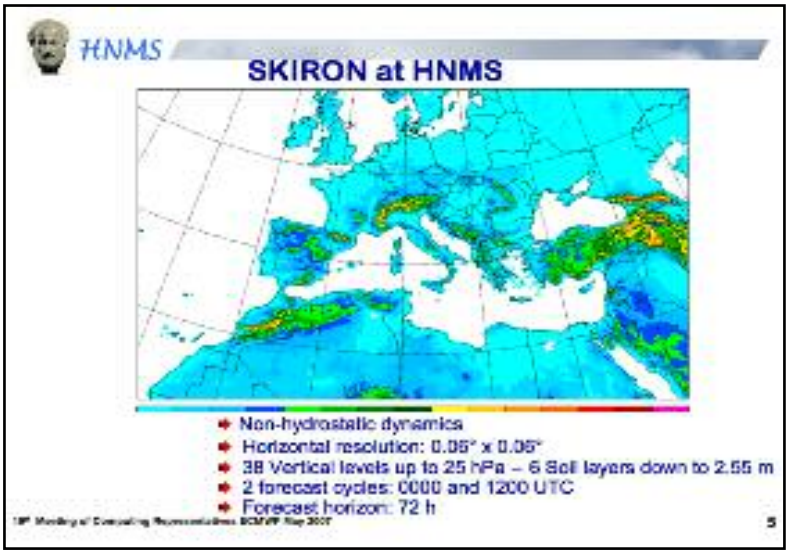

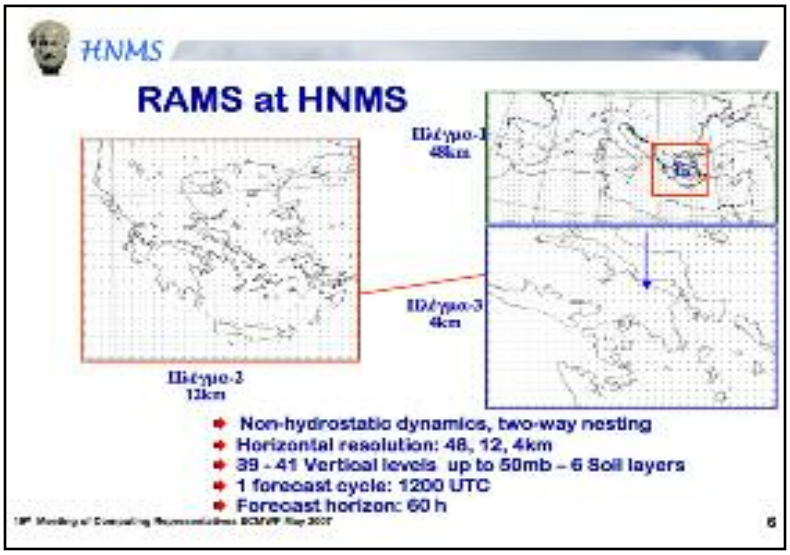

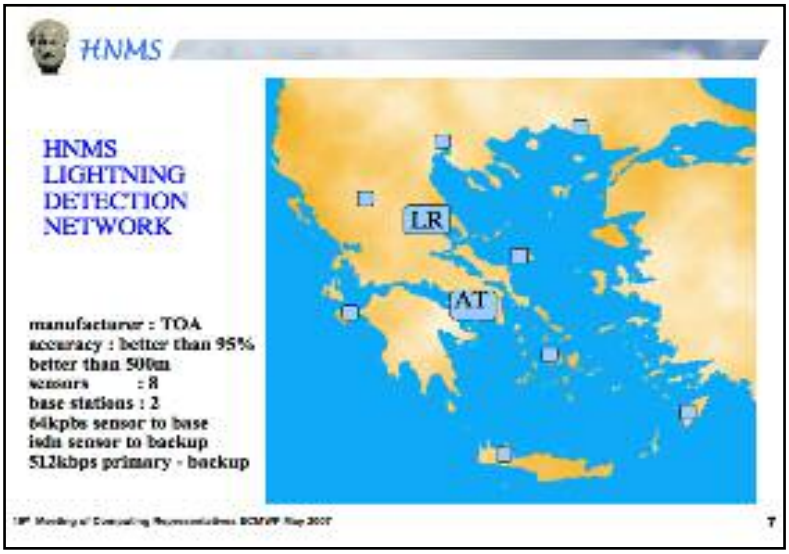

GREECE **GREECE**

R. Urrutia asked what kind of mail server they used. *A. Emmanouil replied it was Postfix and SpamAssassin but noted that SpamAssassin did not work well with Postfix*

I. Weger asked whether operational work was performed on both the HP and IBM systems. *A. Emmanouil replied that the IBM was used solely for operations and the HP for research and backup.* I. Weger then asked whether this meant that they maintained and updated two versions of the model. *A. Emmanouil replied that the versions on the two systems were different.*

HUNGARY

HUNGARY

István Ihász – Hungarian Meteorological Service, Budapest

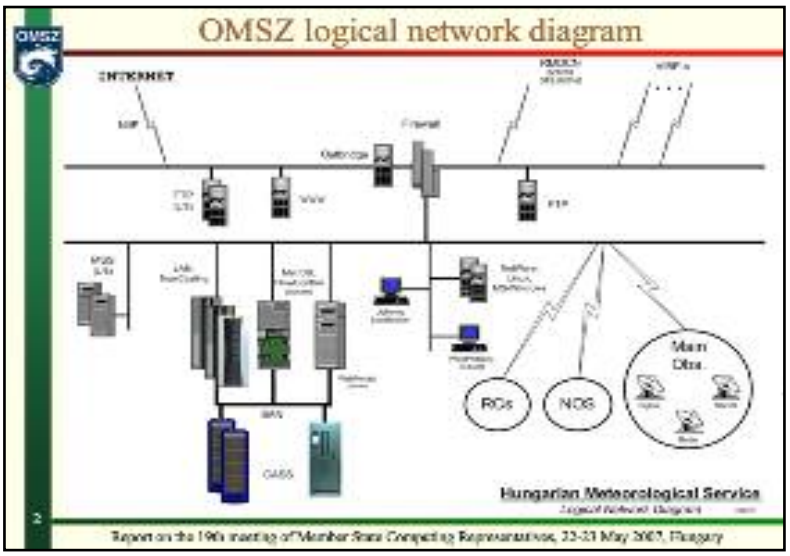

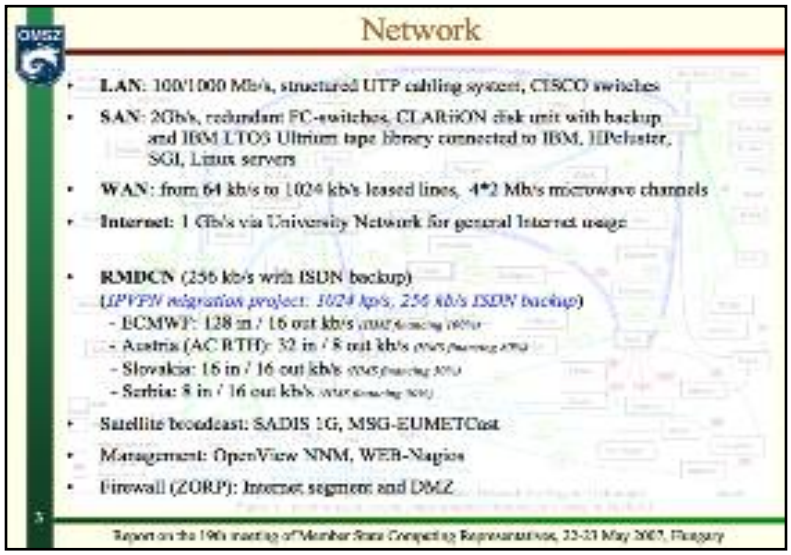

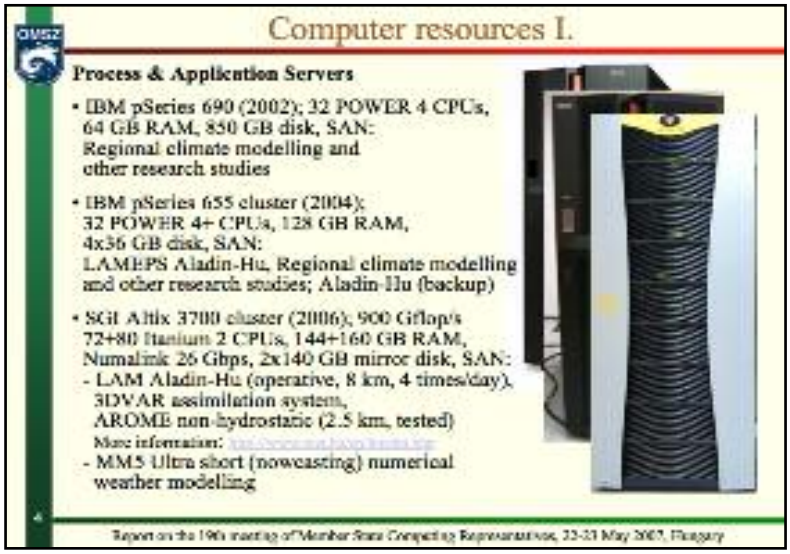

HUNGARY

HUNGARY

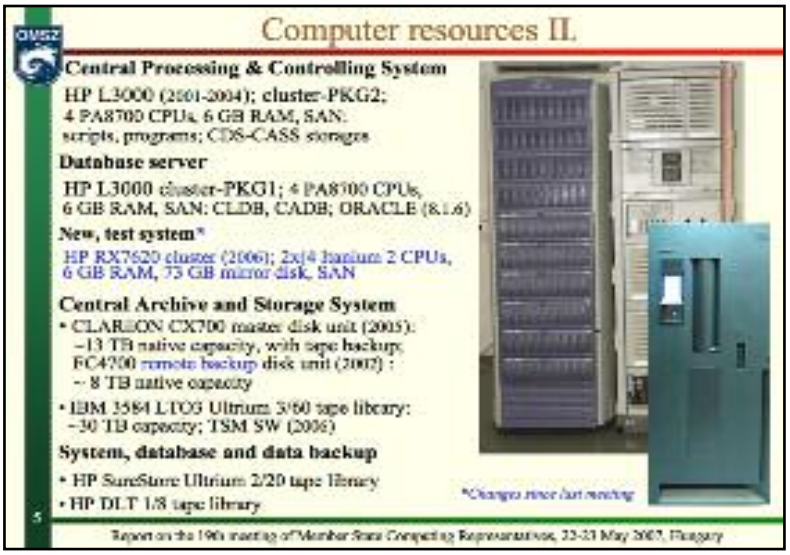

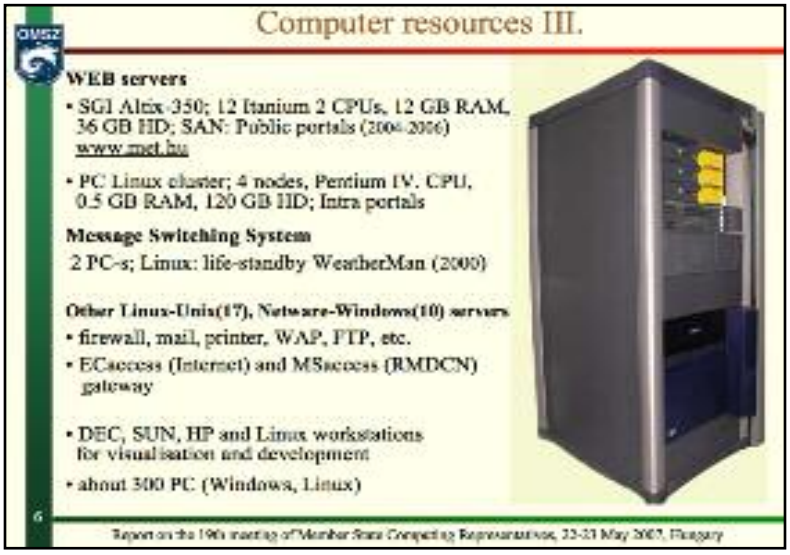

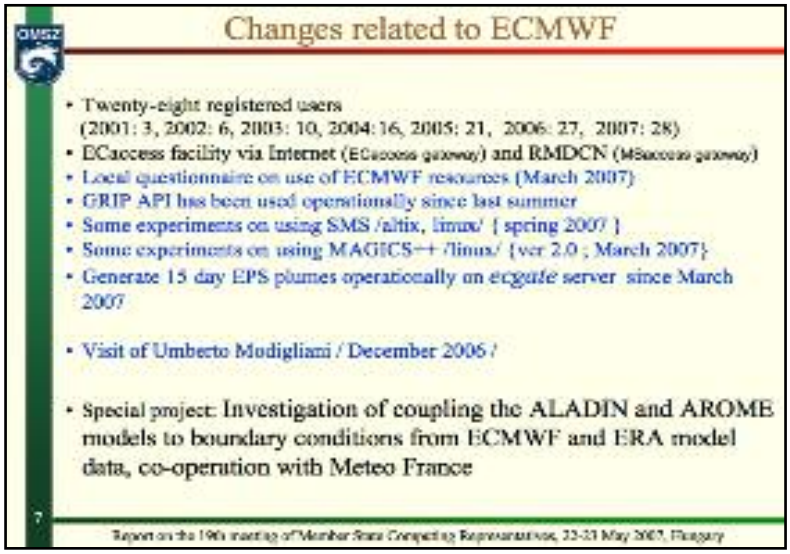

HUNGARY

HUNGARY

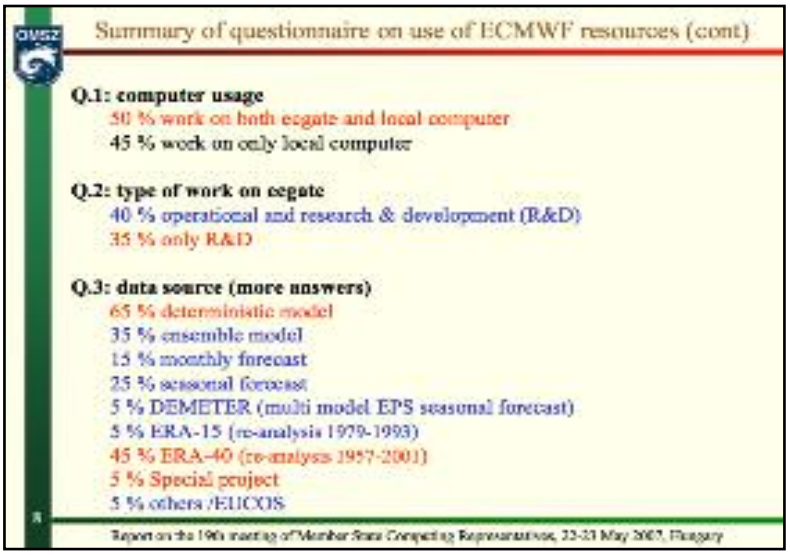

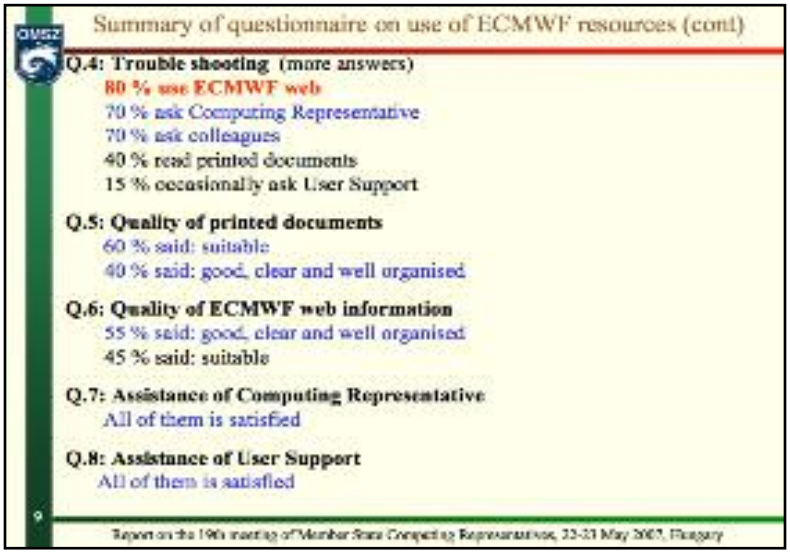

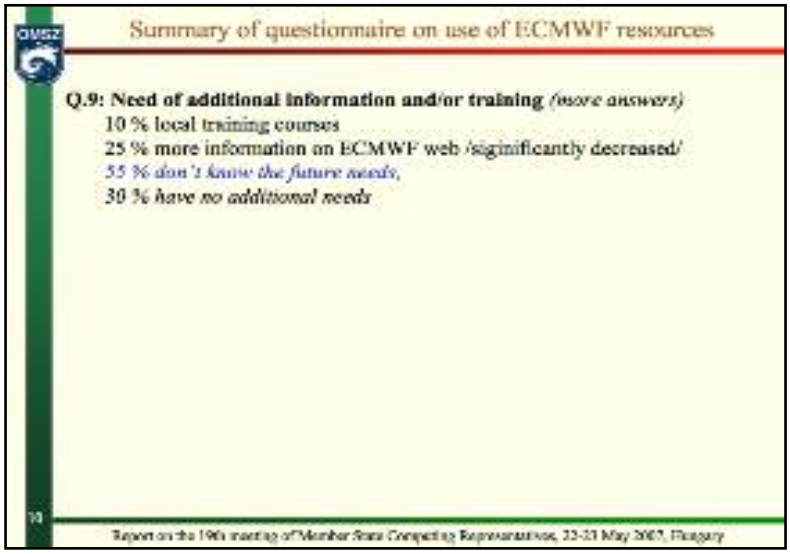

HUNGARY

HUNGARY

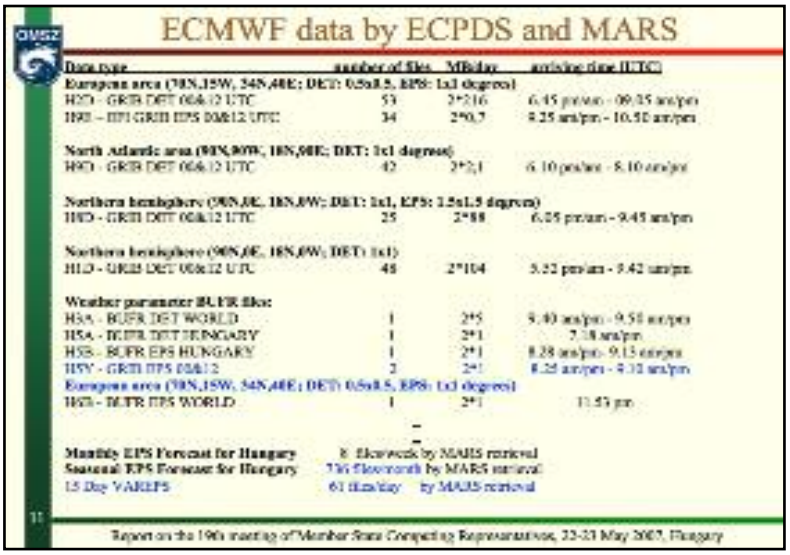

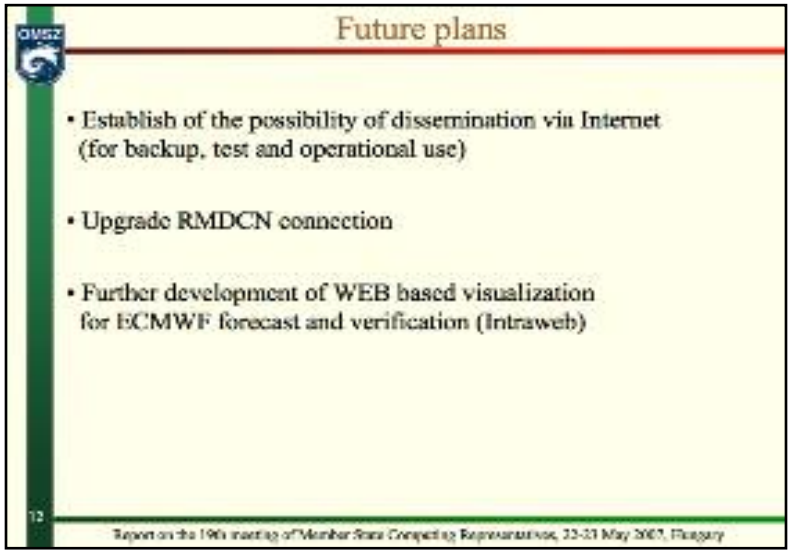

R. Swennen queried their use of SMS with Linux. *I. Ihász replied that ECMWF had provided a special version of SMS for Hungary. R. Swennen said that RMI would like to use SMS on HPUX (Itanium).* U. Modigliani noted that ECMWF already runs SMS on HPUX and could discuss RMI's particular problems.

ICELAND ICELAND

Vigfús Gíslason – Icelandic Meteorological Office

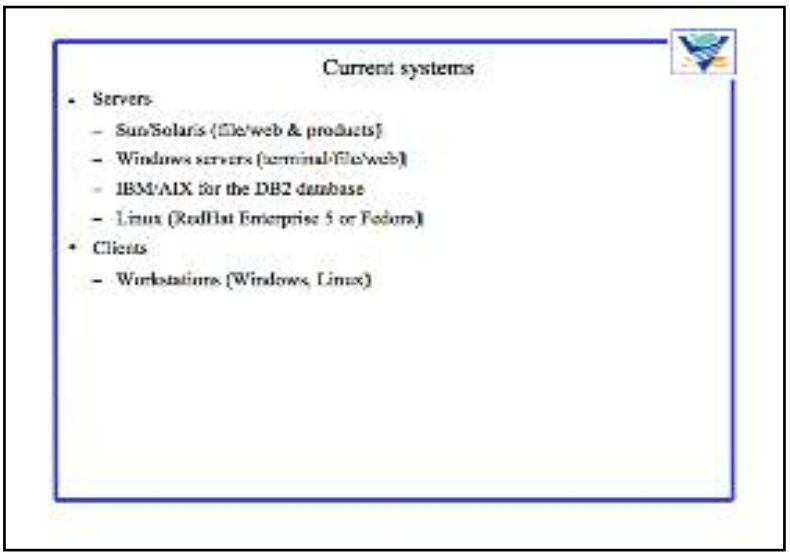

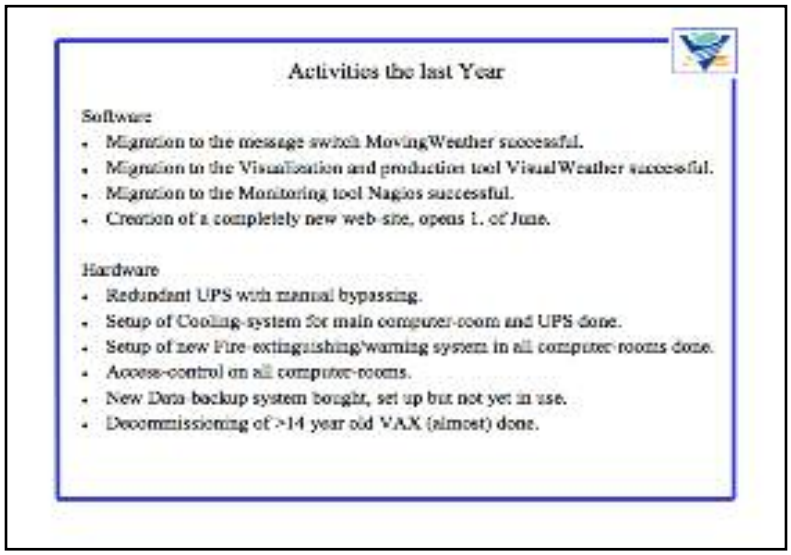

ICELAND ICELAND

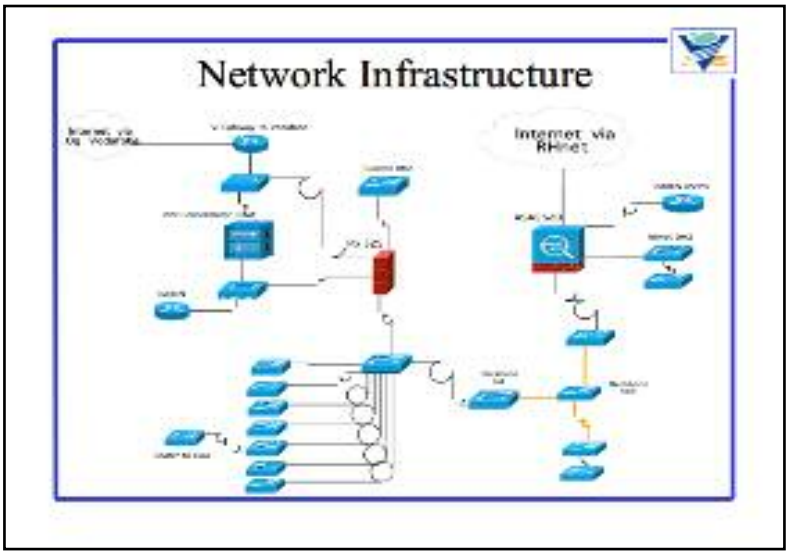

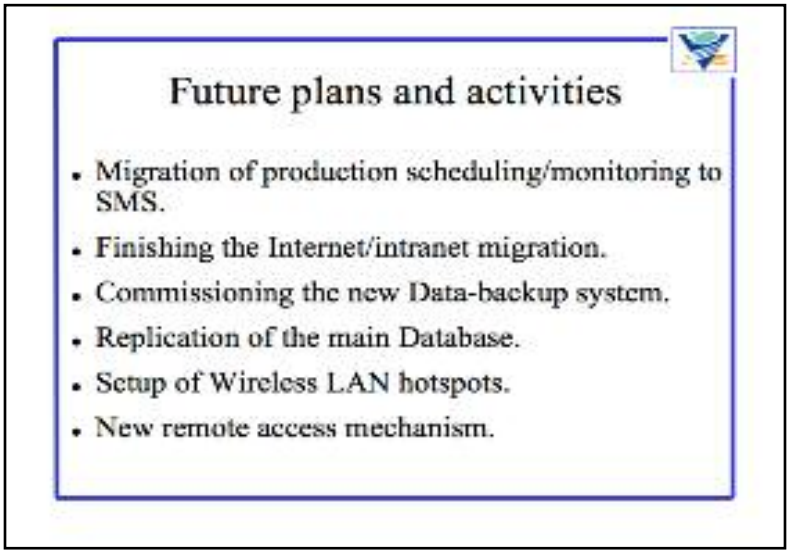

IRELAND

IRELAND

Paul Halton – Met Éireann

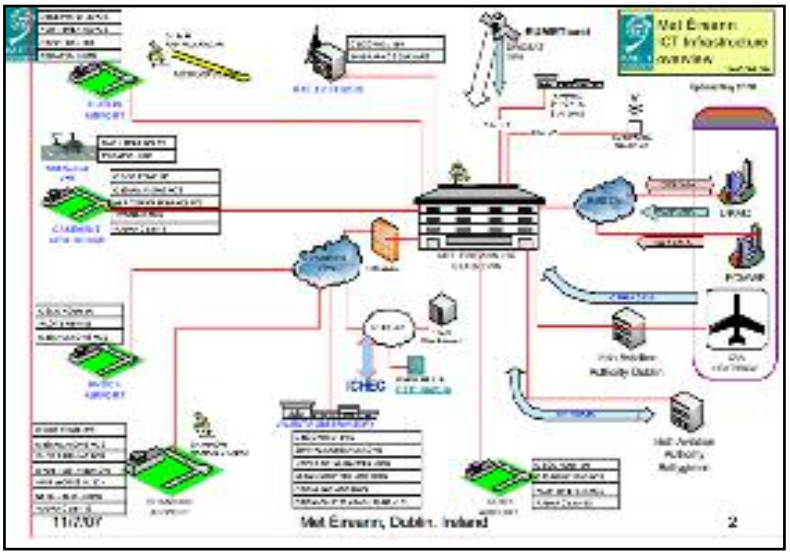

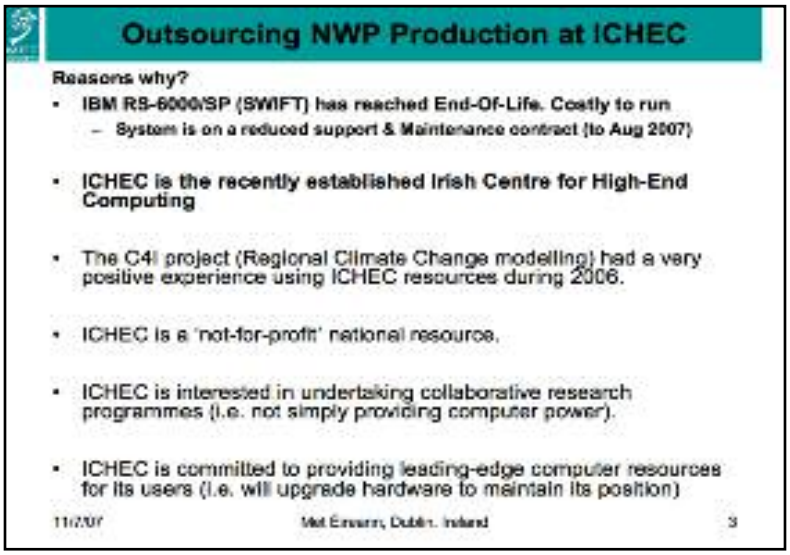

IRELAND IRELAND

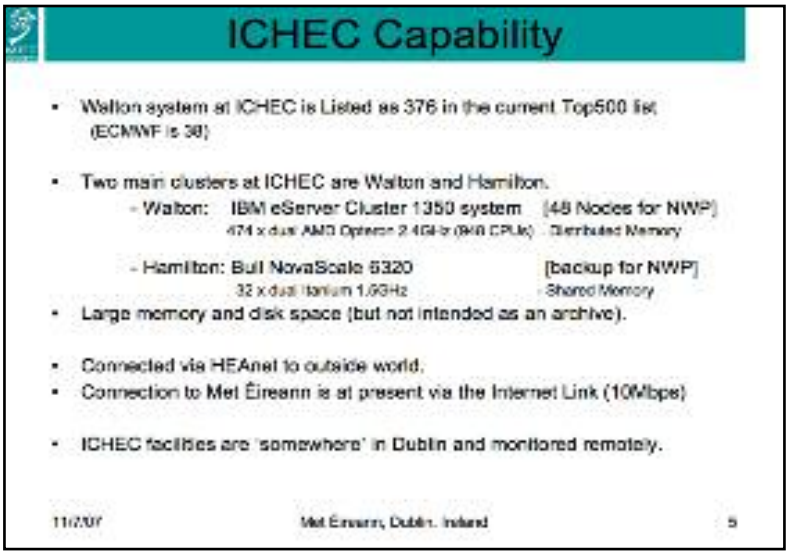

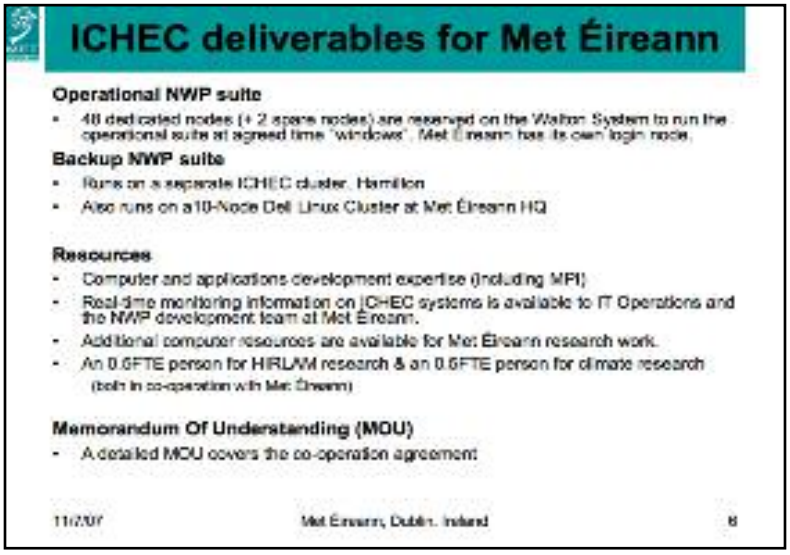

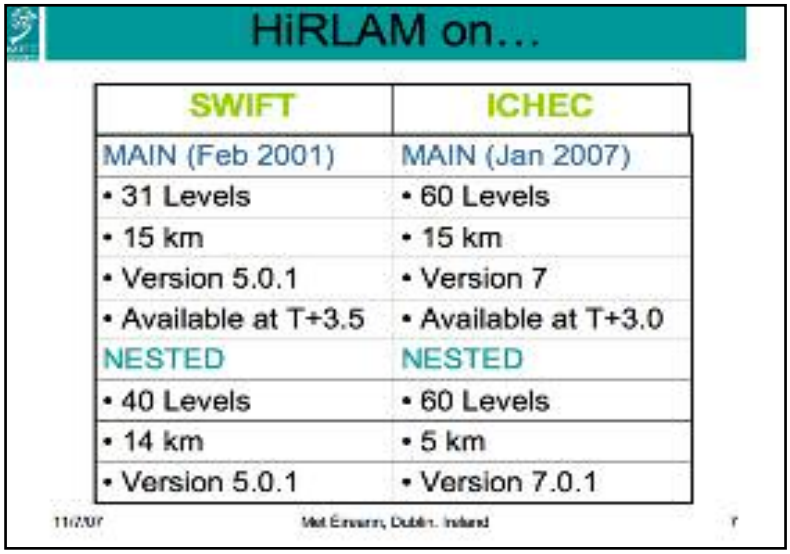

IRELAND

IRELAND

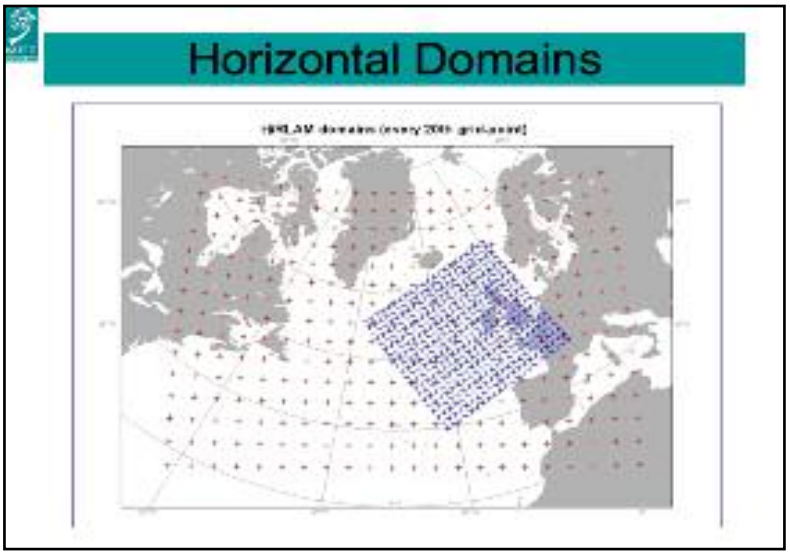

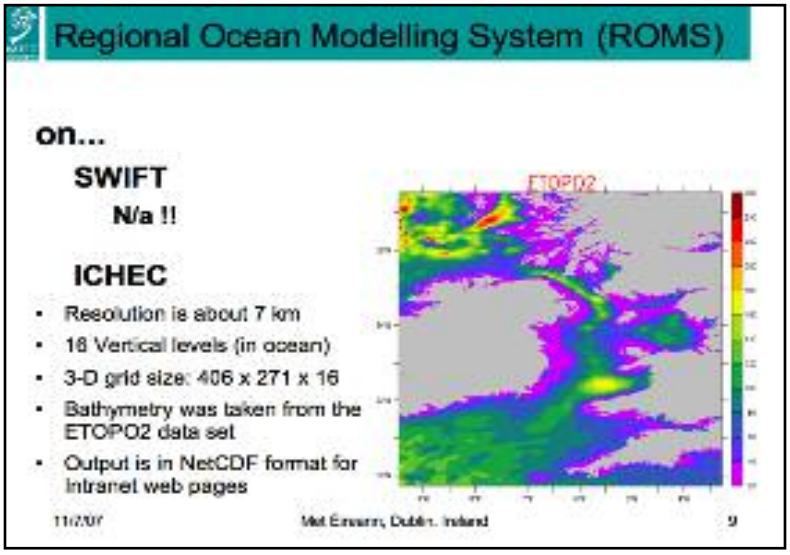

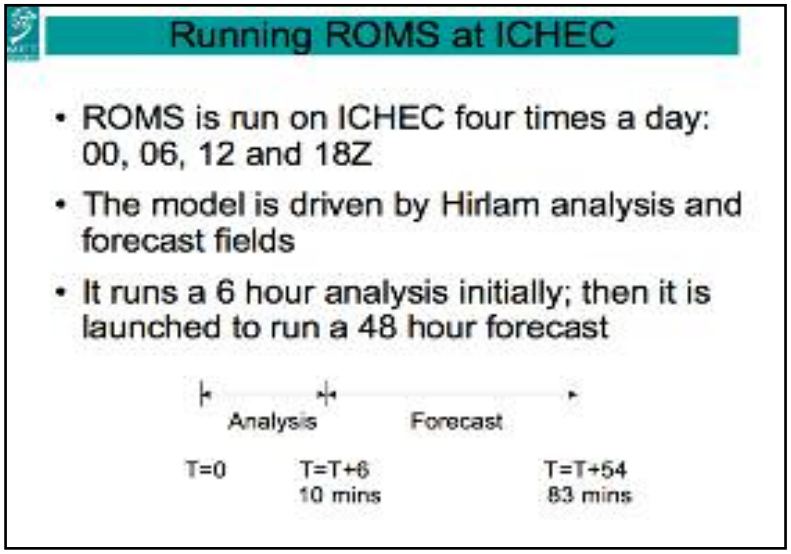

IRELAND

IRELAND

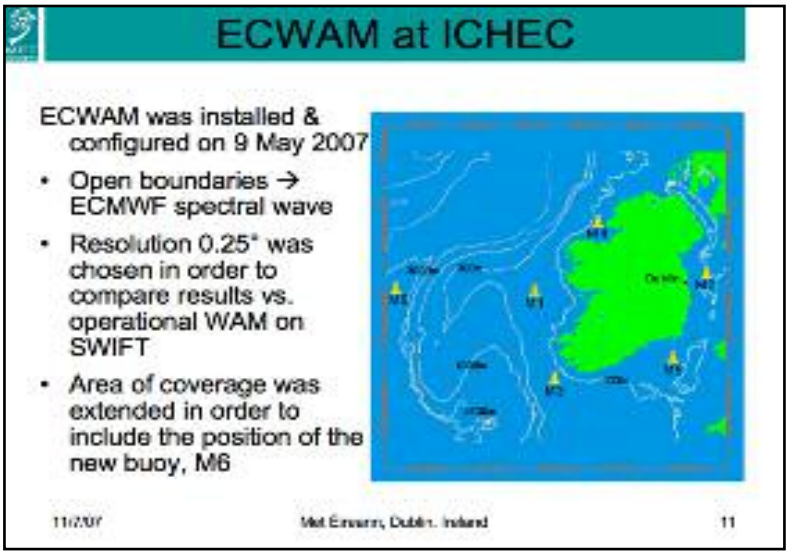

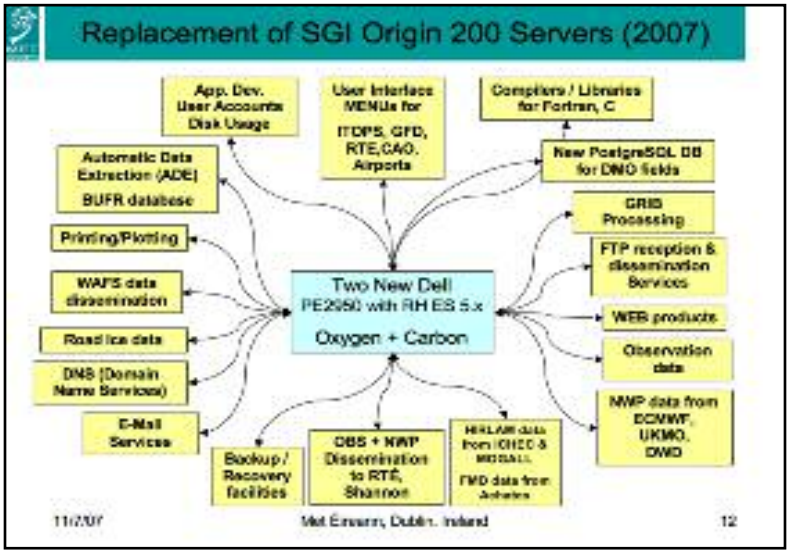

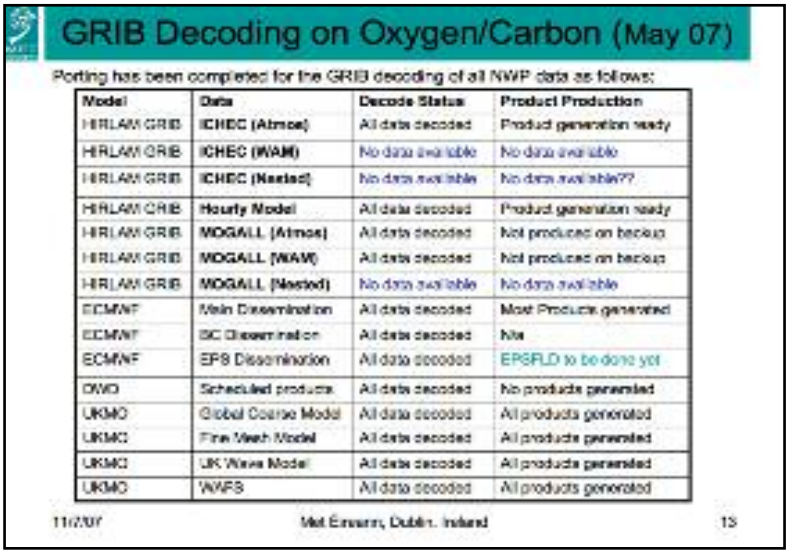

Model	Data	Decode Status	Product Production
HIRLAM GRIB	ICHEC (Atmos)	All data decoded	Product generation ready
HIRLAM GRIB	ICHEC (WAM)	No data available	No data available
HIRLAM GRIB	ICHEC (Nested)	No data available	No data available??
HIRLAM GRIB	Hourly Model	All data decoded	Product generation ready
HIRLAM GRIB	MOGALL (Atmos)	All data decoded	Not produced on backup
HIRLAM GRIB	MOGALL (WAM)	All data decoded	Not produced on backup
HIRLAM GRIB	MOGALL (Nested)	No data available	No data available
ECMWF	Main Dissemination	All data decoded	Most Products generated
ECMWF	BC Dissemination	All data decoded	N/a
ECMWF	EPS Dissemination	All data decoded	EPSRUD to be done yet
DWD	Scheduled products	All data decoded	No products generated
UKMO	Global Coarse Model	All data decoded	All products generated
UKMO	Fine Mesh Model	All data decoded	All products generated
UKMO	UK Wave Model	All data decoded	All products generated
UKMO	WAFS	All data decoded	All products generated

IRELAND

IRELAND

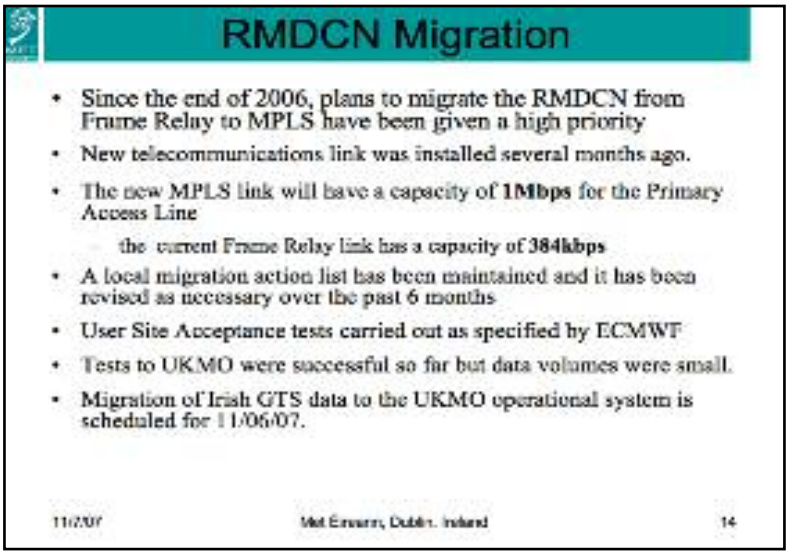

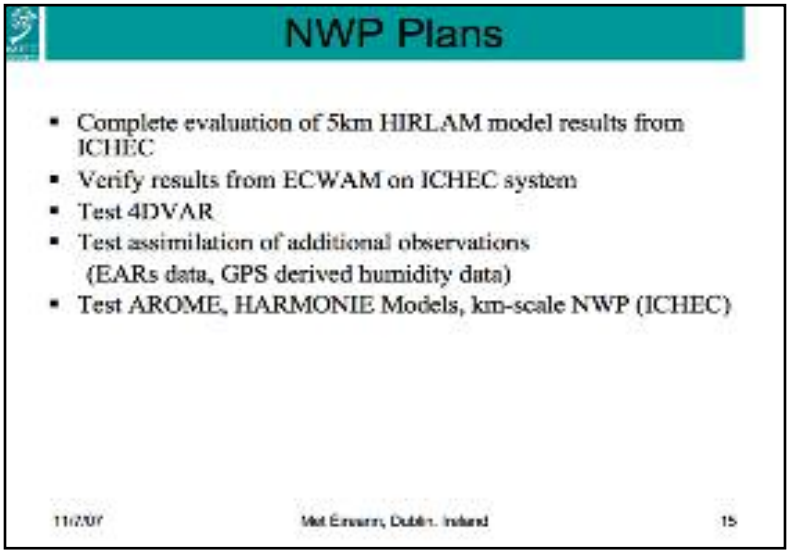

IRELAND IRELAND

Special Projects

SPIE C4I Project

- C4I Project is a major user of the ECMWF resources at this stage.
- The project in its current format is now entering its final months.
- Climate Change research in Ireland is expected to obtain more permanent funding after the next Government is formed in summer 2007.
- We expect that new projects will emerge towards the end of 2007.

SPIE SERG Project

- This project has been completed and a final report is expected soon.
- **The researchers at UCC are still seeking funding for a new project**

Boundary Conditions (BC) Optional Project

- Since July 2006, Met Éireann has been receiving all BC vertical levels in the range: surface level -10hpa level [i.e. levels 18-90]
- The extra levels are included in the routine dissemination of the BC data used with the operational HIRLAM model runs on SWIFT and ICHEC platforms.

11/2/07 Met Éireann, Dublin, Ireland 17

ActivIdentity Tokens and SecurID Cards

34 Member State Users with Tokens / Cards Enabled

- 20 Users received new ActivIdentity Tokens in April 2007
- 2 Users at UCC did not receive their Tokens in the post
- 2 New Users were assigned new ActivIdentity Tokens on 17 May
- 10 Users continue to use their valid and working SecurID cards

The Expired SecurID cards were returned to Computing Representative

Currently we have:

- 3 Spare ActivIdentity Tokens
- 2 Spare SecurID Cards

Registration of user accounts is very convenient and efficient via the ECMWF web pages

11/2/07 Met Éireann, Dublin, Ireland 18

ECaccess and new Tokens

To use ECaccess with the new tokens, the ECaccess client software had to be updated on PCs /servers. The new version is v3.1.0.

Users were required to follow these Steps:

1. Download the ECaccess client from ECMWF
   http://www.ecmwf.int/services/ecaccess/download/ecaccess-client-v3.1.0.tar.gz
2. Install new client software either in own account or under /usr/local
   (In /usr/local, change permissions using chmod -R o+rx .)
3. Edit **ecaccess-v3.1.0/client/tools/.ecaccess**
   1. Insert the following lines:-
      - _ECACCESS_HOME=/usr/local/ecaccess-v3.1.0 (user's local path)
      - _ECHOST=123.456.789.111 (IP address of local server)
      - _ECFTPPORT=9021
      - _ECCERTPORT=9443
4. Edit .bashrc or .cshrc file changing ecaccess-v2.0.3 to ecaccess-v3.1.0

11/2/07 Met Éireann, Dublin, Ireland 19

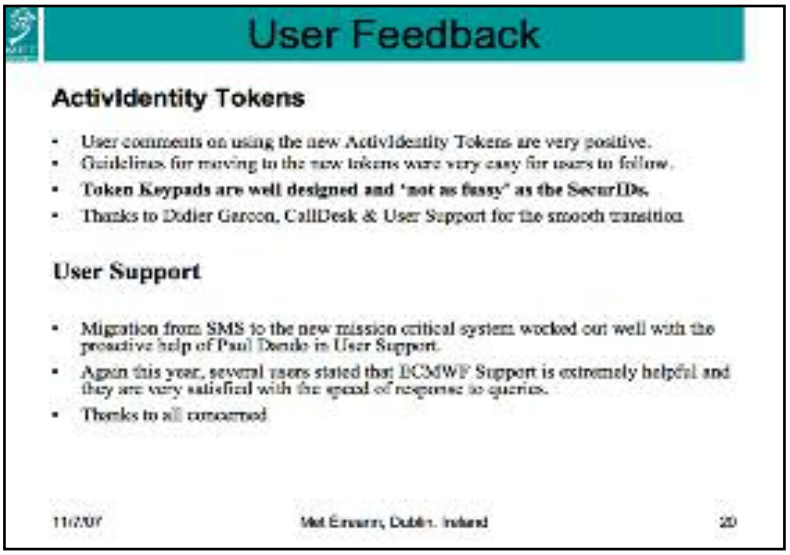
User Feedback
ActivIdentity Tokens
User comments on using the new ActivIdentity Tokens are very positive.
Guidelines for moving to the new tokens were very easy for users to follow.
Token Keypads are well designed and 'not as fussy' as the SecurIDs.
Thanks to Didier Garcon, CallDesk & User Support for the smooth transition
User Support
Migration from SMS to the new mission critical system worked out well with the proactive help of Paul Dando in User Support.
Again this year, several users stated that ECMWF Support is extremely helpful and they are very satisfied with the speed of response to queries.
Thanks to all concerned
11/7/07
Met Éireann, Dublin, Ireland
20

ITALY ITALY

Carmelo Gambuzza – CNMCA, Italian Met Service

Last year's changes at the Italian Met. Service:

New Mail Server
Upgrade of the Internet link bandwidth
Completed the national radar network setup
Completed the national atmospheric lightning survey network setup

Use of ECMWF facilities

Upgrade of the Internet bandwidth:

Old: 2 Mbsp
New: 10 Mbps
Before the end of 2007: 100 Mbps

New Mail Server

Old
- HW: 1 HP Proliant ML350 G2 x86
- OS: RH Linux

New
- HW: 2 IBM eServer 325 Type8835 (2x Opteron 64bit 2.2 Ghz + 2 Gbyte)
- Raid1
- OS: RedHat Linux Enterprise ES3 - update 6
- Cluster HA: HeartBeat + DRBD (latest stable version)

ECMWF, 22-24 May 2007

Meteorological RADAR Network

GPM 500 C

Connected with 384 Kbps leased links

RADAR Network

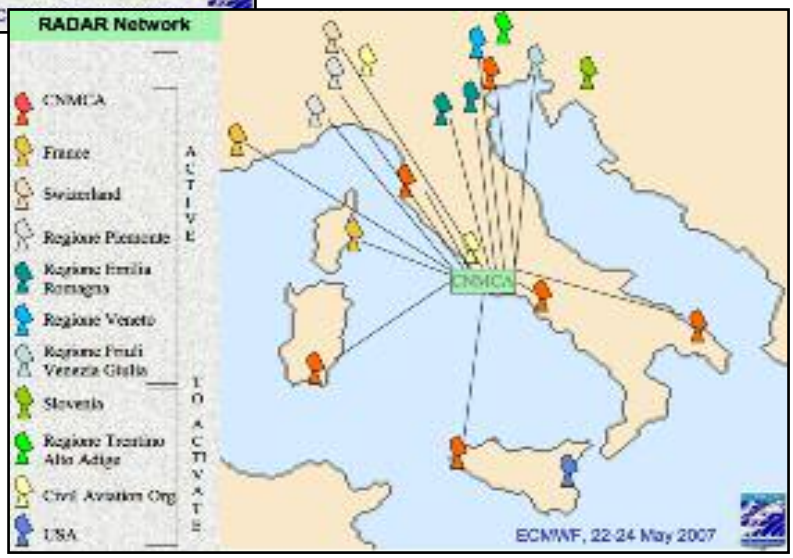

ECMWF, 22-24 May 2007

ITALY

ITALY

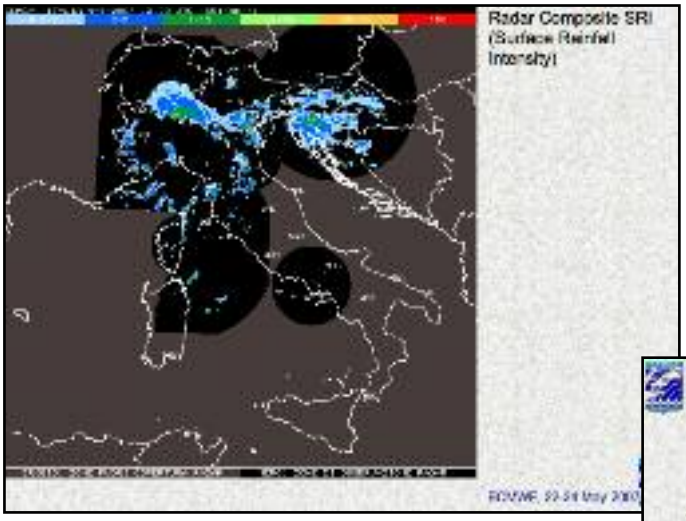
Radar Composite SRI
(Surface Rainfall
Intensity)

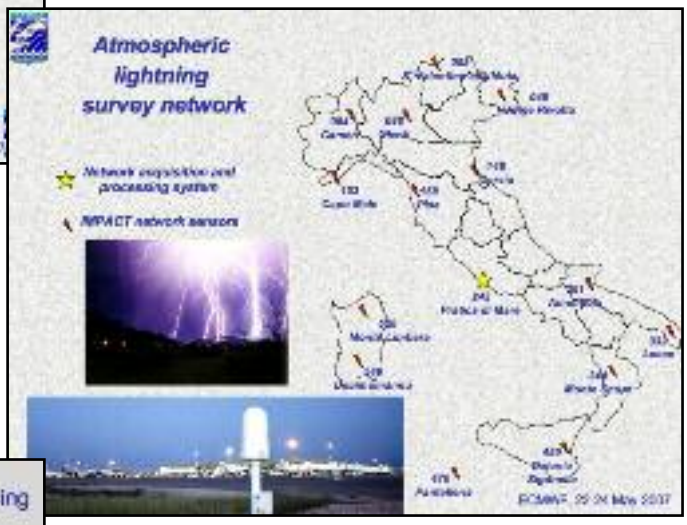
Atmospheric
lightning
survey network


LampiNet
(Atmospheric lightning
network)


Lampinet

ITALY

ITALY

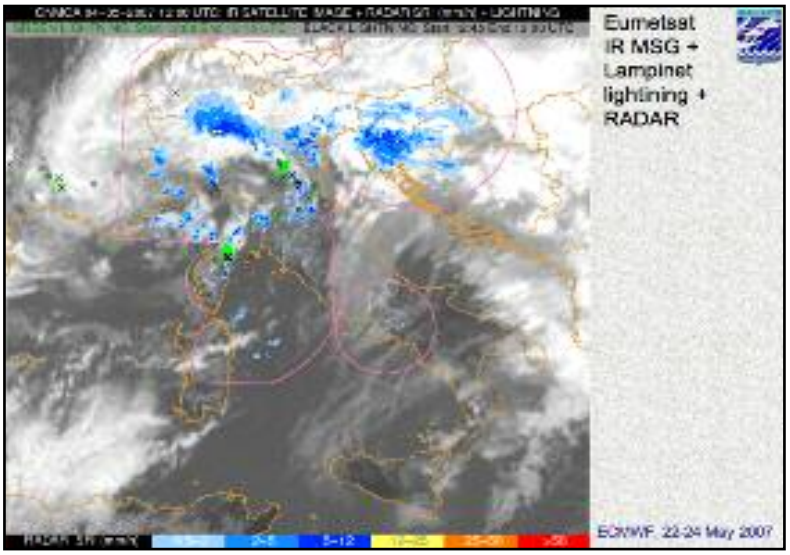

Use of ECMWF facilities

ECMWF Products:

Total of 1.5 Gbytes of dissemination GRIB data received per day:

- 700 Mbytes for internal purposes
- 800 Mbytes for external users:
  - ✓ Regional Met. Services
  - ✓ Universities
  - ✓ Research Centres
  - ✓ Civil Protection Organization
  - ✓ Antartic research programme
  - ✓ Italian Space Agency
  - ✓ European Space Agency
  - ✓ TV Networks
  - ✓ Motorways society
  - ✓ Electrical Power company
  - ✓ Other public and private companies

ECMWF, 22-24 May 2007

Use of ECMWF facilities

ECMWF, 22-24 May 2007

ECMWF Users & Projects

Total of 150 users:
- 50 from National Met Service
- 100 from other institutions

✓ Retrieval of MARS data (mainly for external users)

✓ Two time critical applications on HPCF:
- Italian Met Service NWP System
  - 3DVAR assimilation cycle (14 km)
  - COSMO-ME (7 km)
  - COSMO-IT (2.8 km)
- COSMO – LEPS (10 km)

✓ Experimental NWP (EnKF, etc) suites on HPCF

✓ Ecaccess 3.1.0 installed on 8-3-2007 (thanks to Paul Dando & Laurent Gougeon)

✓ Use of Actividentity tokens without any problem

Total of 12 Special Projects

MOROCCO

MOROCCO

Hassan Haddouch – Moroccan Meteorological Centre

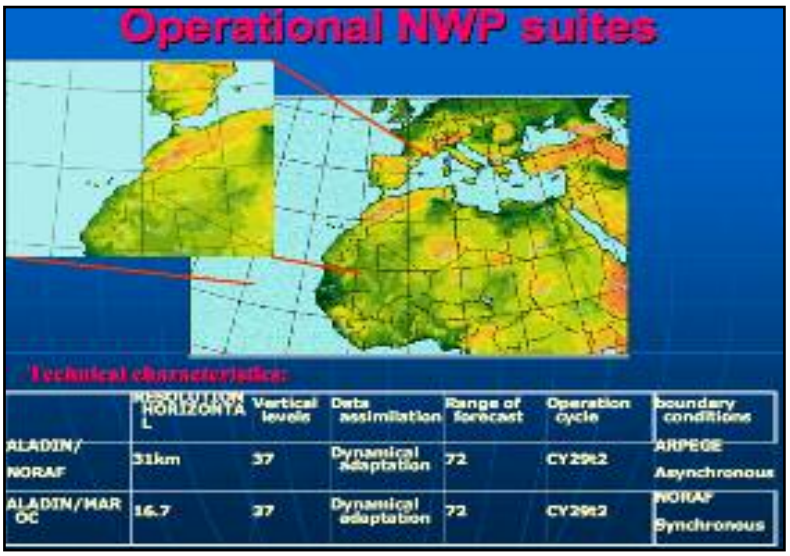

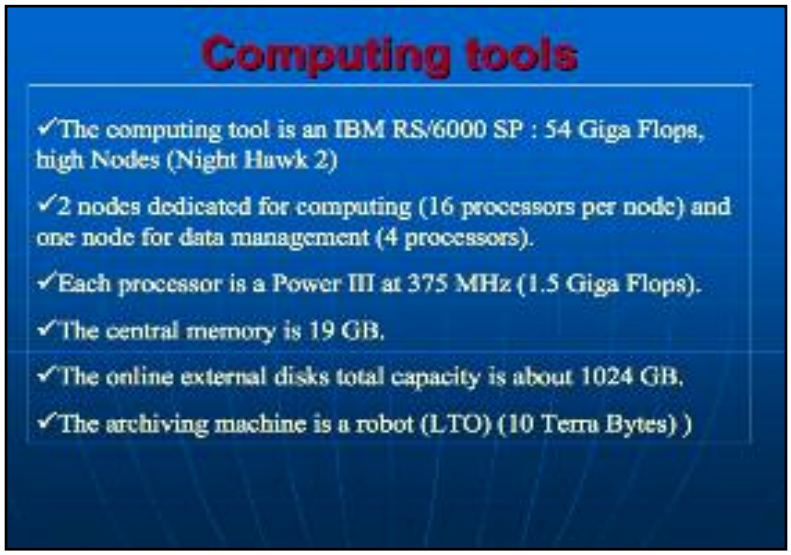

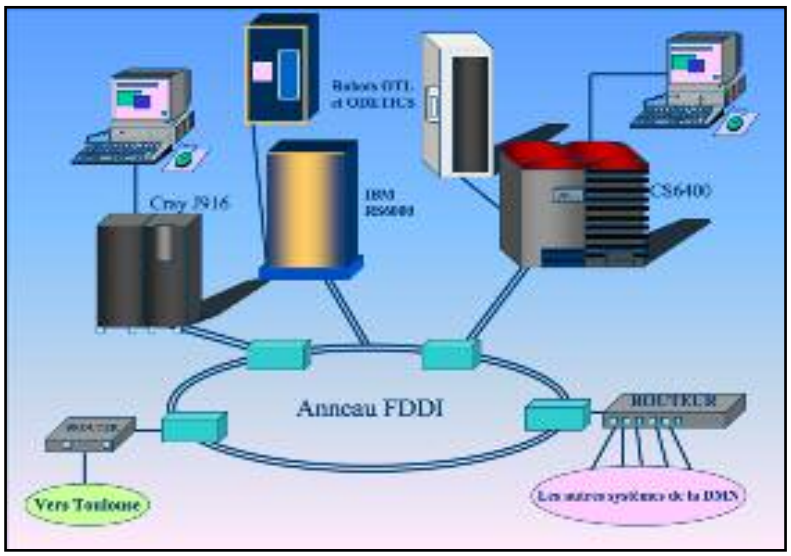

MOROCCO MOROCCO

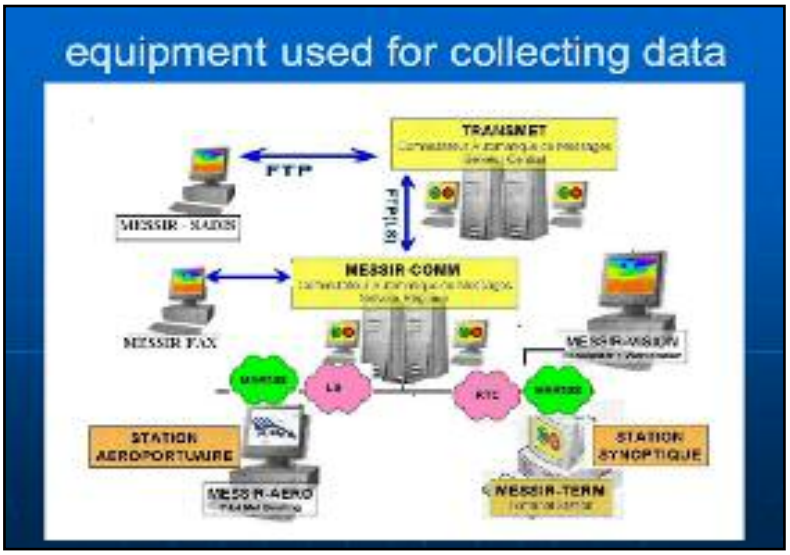

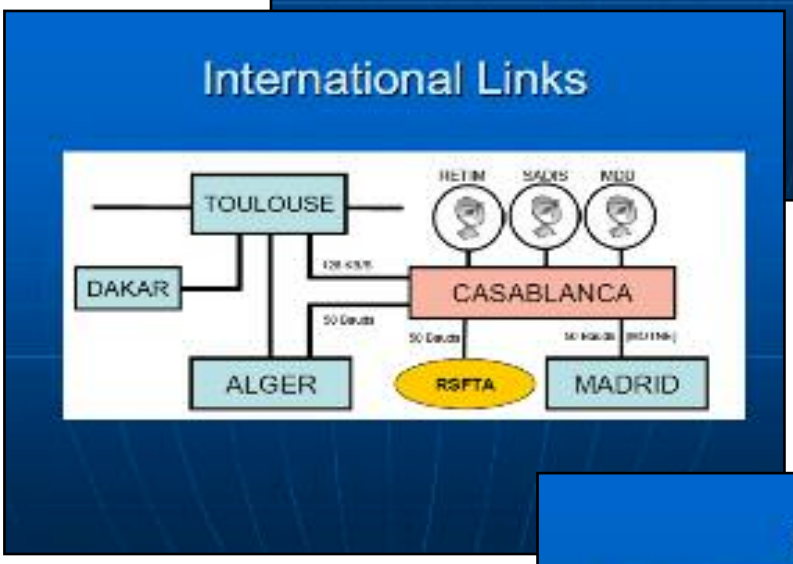

NETHERLANDS

NETHERLANDS

Hans de Vries – KNMI

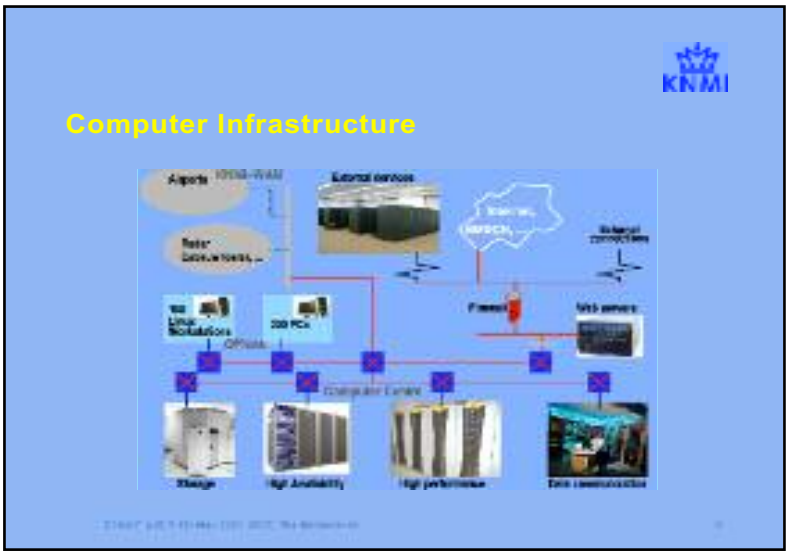

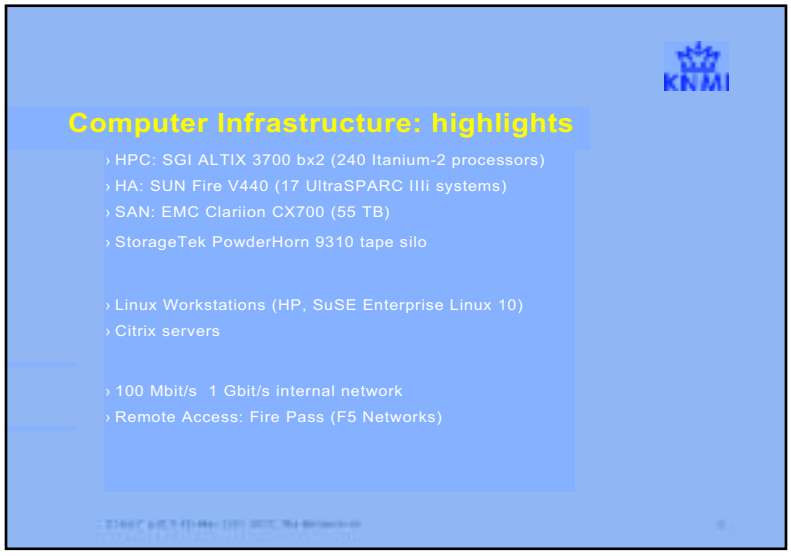

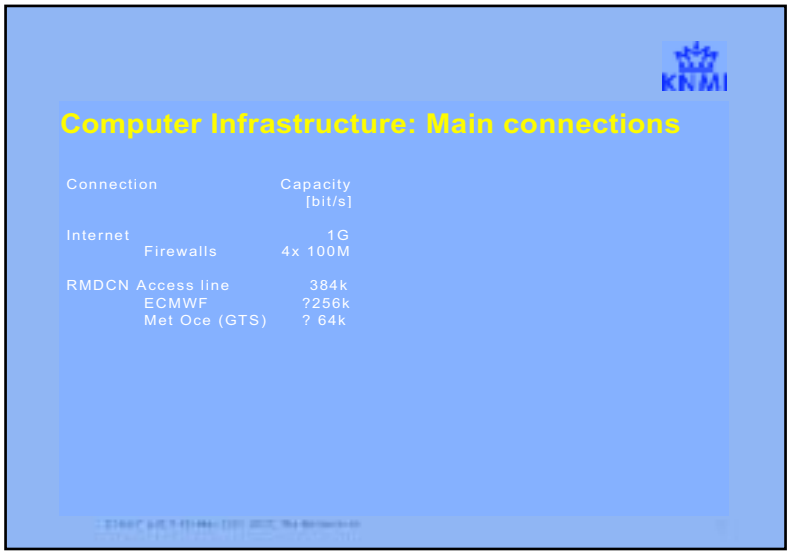

NETHERLANDS

NETHERLANDS

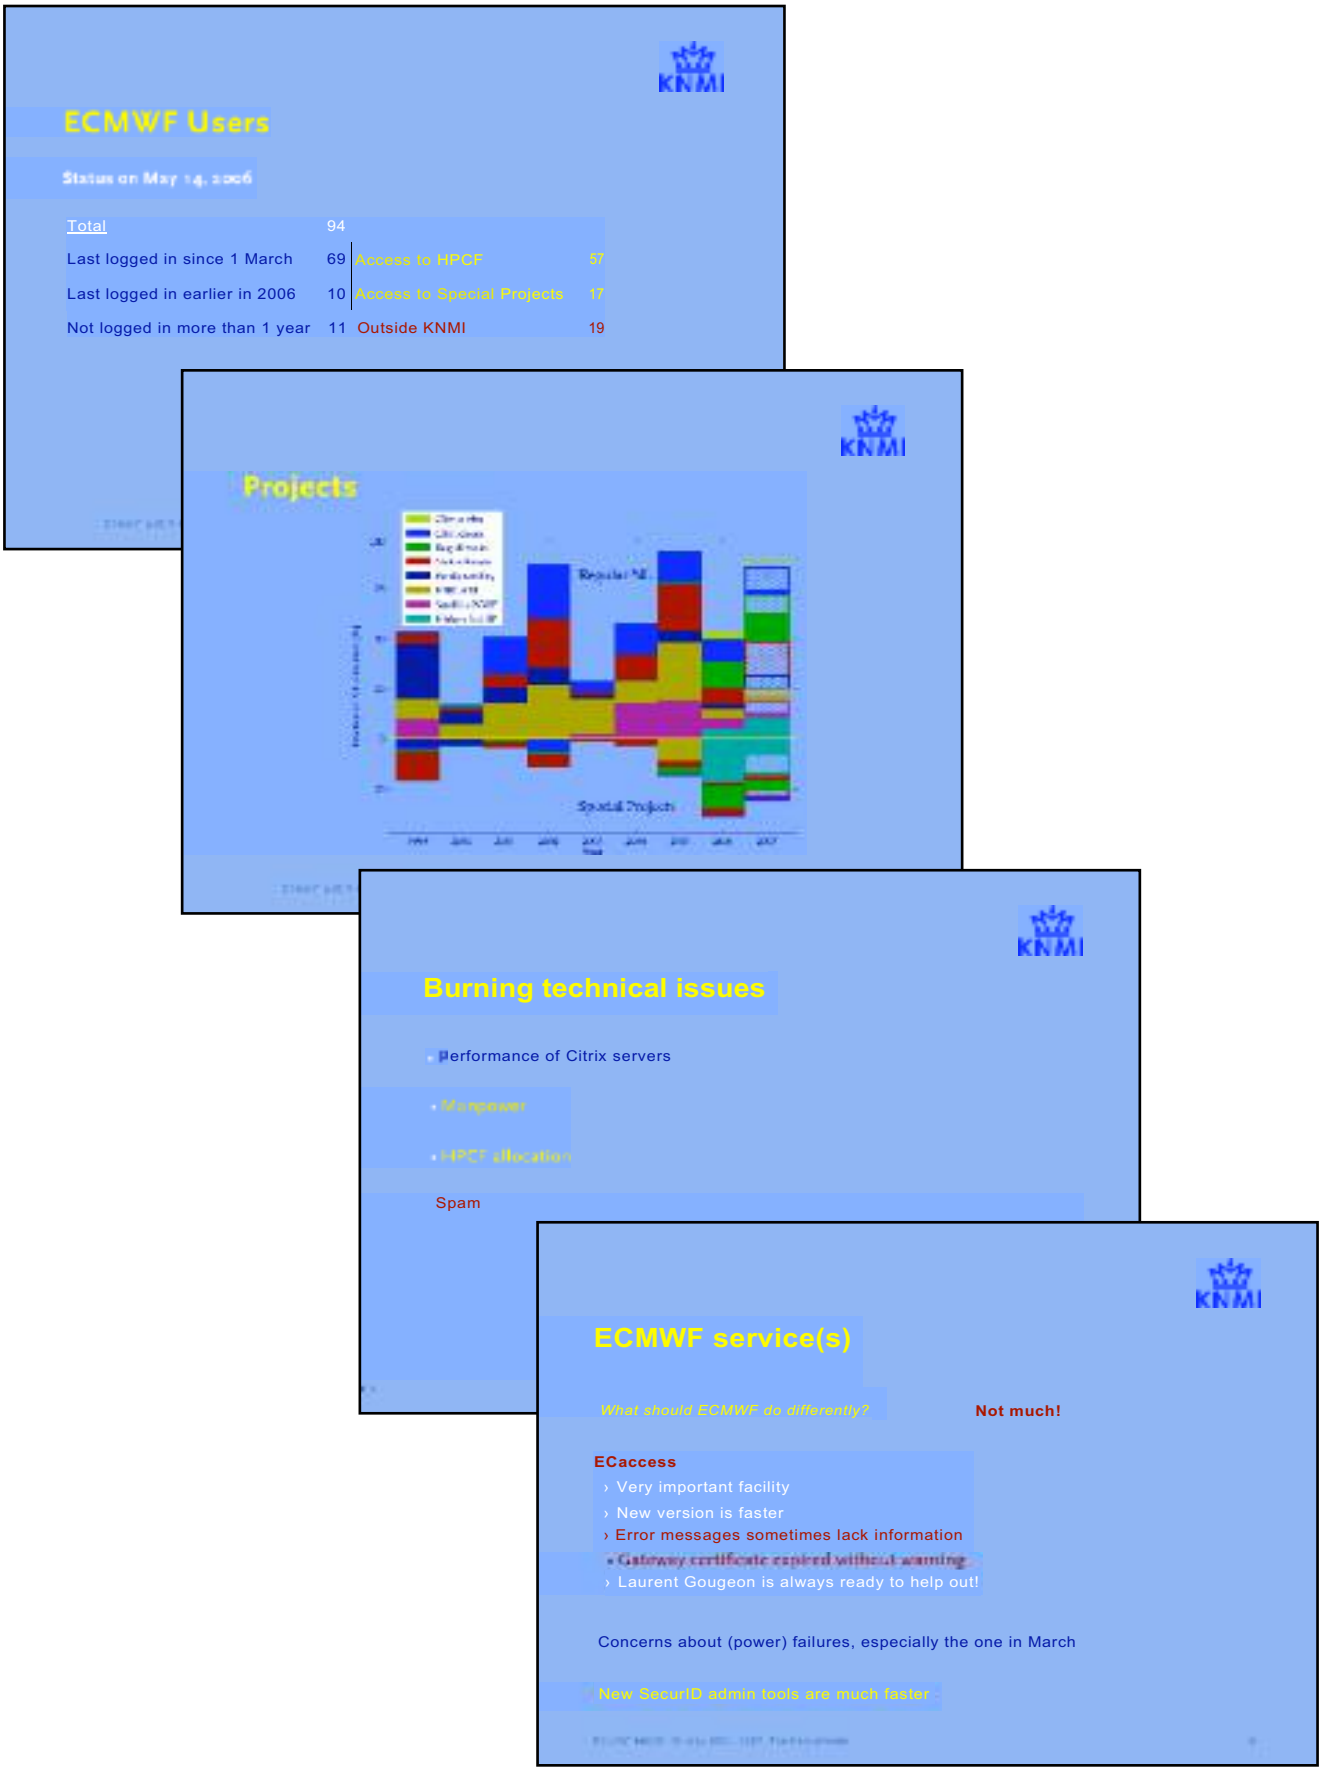

NORWAY NORWAY

Rebecca Rudsar – Norwegian Meteorological Institute

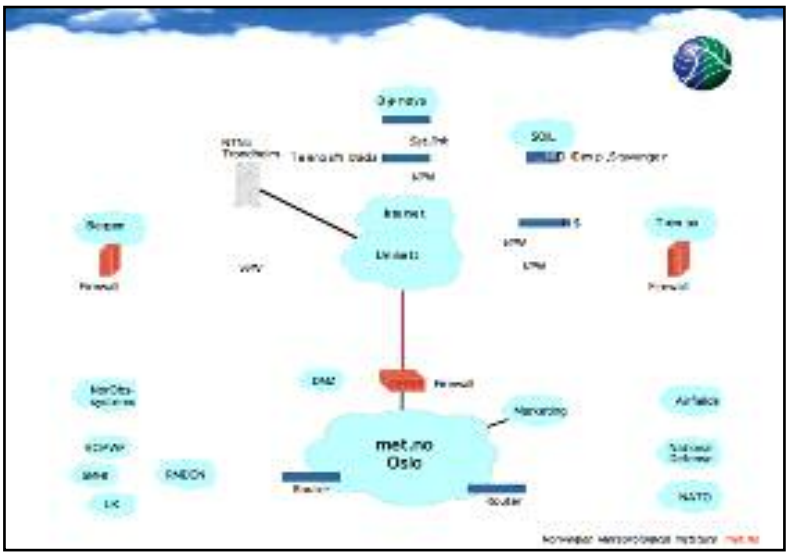

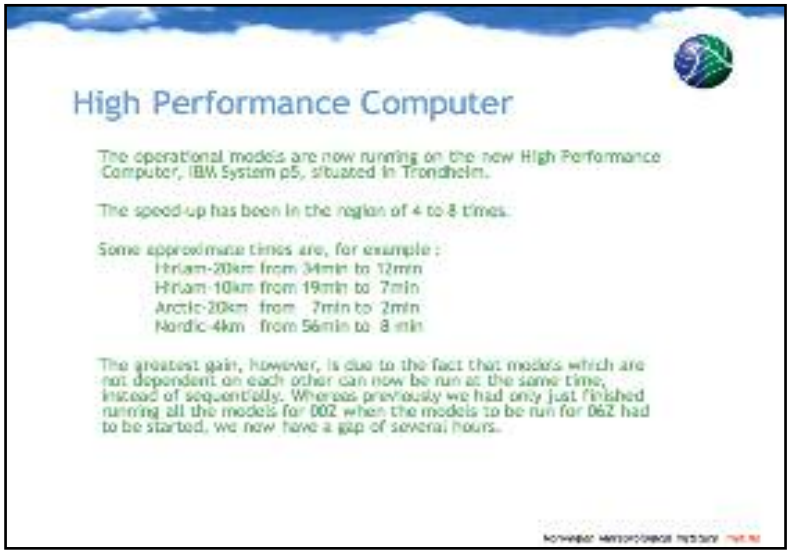

NORWAY

NORWAY

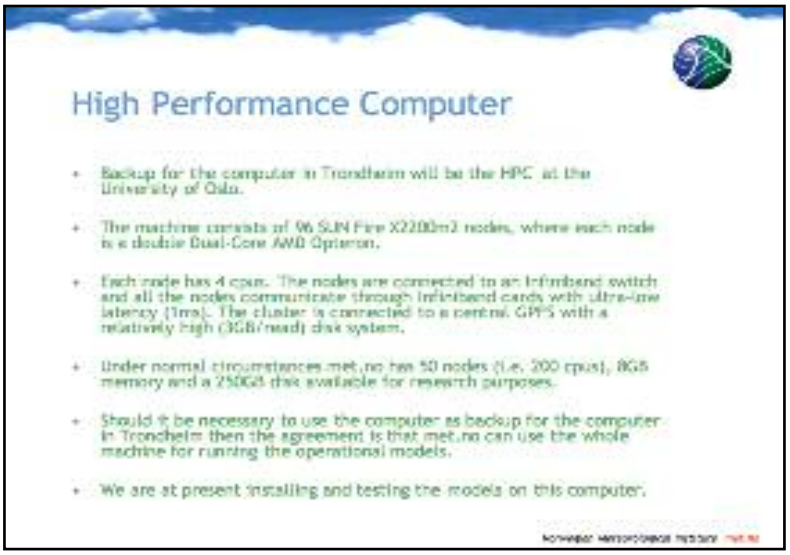

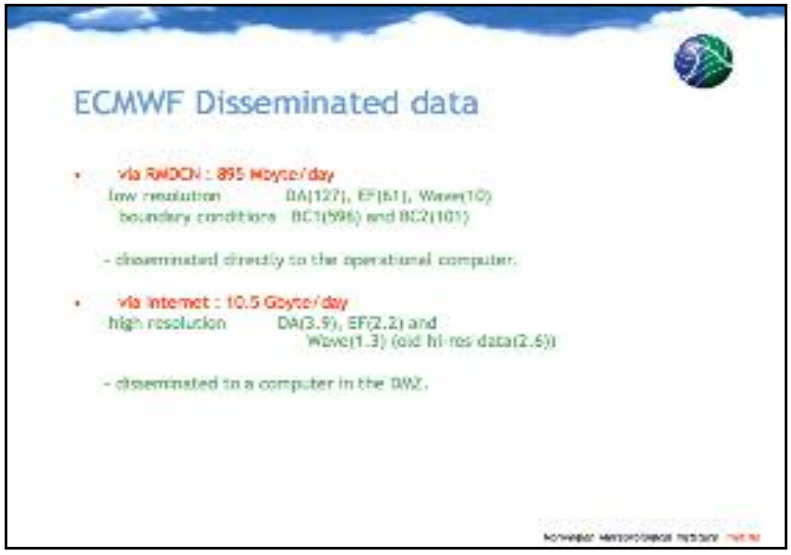

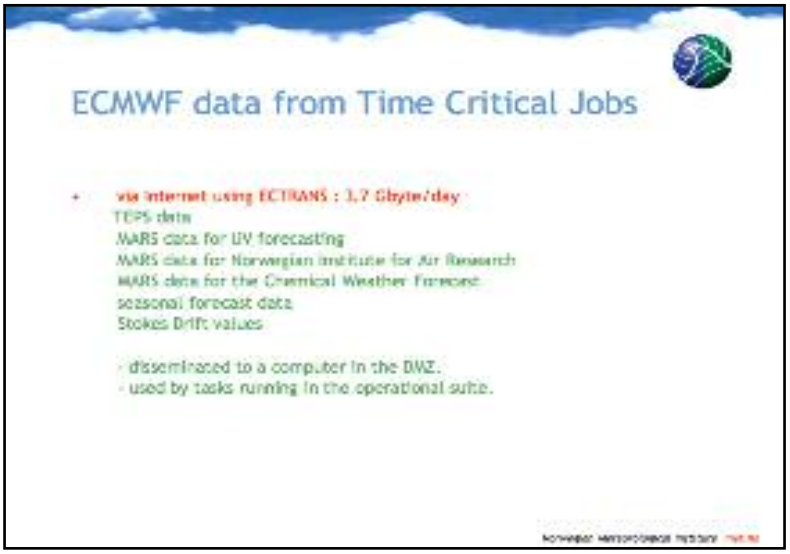

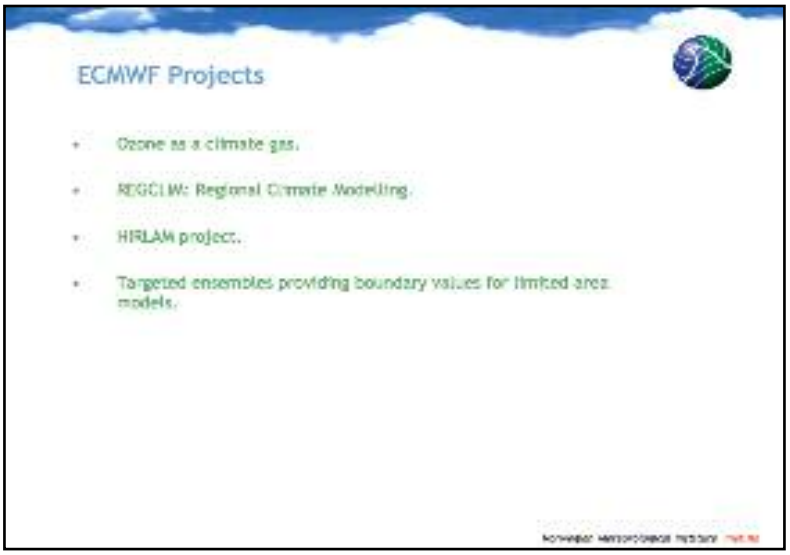
ECMWF Projects
Ozone as a climate gas.
REGCLIM: Regional Climate Modelling.
HIRLAM project.
Targeted ensembles providing boundary values for limited area models.

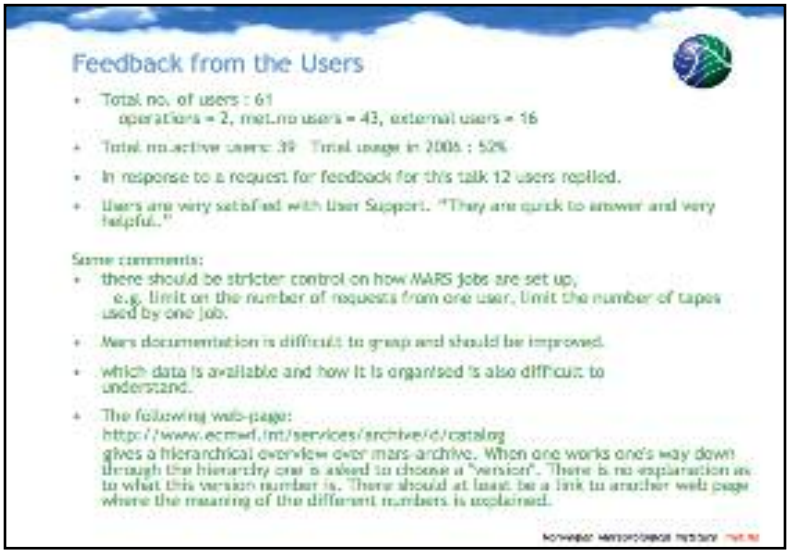
Feedback from the Users
Total no. of users : 61
operations = 2, met.no users = 43, external users = 16
Total no.active users: 39 Total usage in 2006 : 52%
In response to a request for feedback for this talk 12 users replied.
Users are very satisfied with User Support. "They are quick to answer and very helpful.."
Some comments:
there should be stricter control on how MARS jobs are set up,
e.g. limit on the number of requests from one user, limit the number of tapes used by one job.
Mars documentation is difficult to grasp and should be improved.
which data is available and how it is organised is also difficult to understand.
The following web-page:
http://www.ecmwf.int/services/archive/d/catalog
gives a hierarchical overview over mars archive. When one works one's way down through the hierarchy one is asked to choose a "version". There is no explanation as to what this version number is. There should at least be a link to another web page where the meaning of the different numbers is explained.

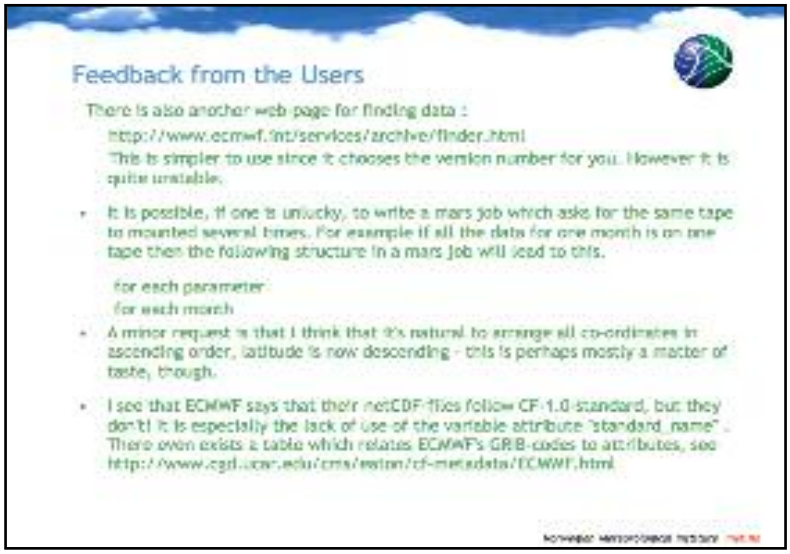
Feedback from the Users
There is also another web page for finding data :
http://www.ecmwf.int/services/archive/finder.html
This is simpler to use since it chooses the version number for you. However it is quite unstable.
It is possible, if one is unlucky, to write a mars job which asks for the same tape to mounted several times. For example if all the data for one month is on one tape then the following structure in a mars job will lead to this.
for each parameter
for each month
A minor request is that I think that it's natural to arrange all co-ordinates in ascending order, latitude is now descending - this is perhaps mostly a matter of taste, though.
I see that ECMWF says that their netCDF files follow CF-1.0 standard, but they don't! It is especially the lack of use of the variable attribute "standard_name". There even exists a table which relates ECMWF's GRIB-codes to attributes, see http://www.cgd.ucar.edu/cms/eaton/cf-metadata/ECMWF.html

NORWAY NORWAY

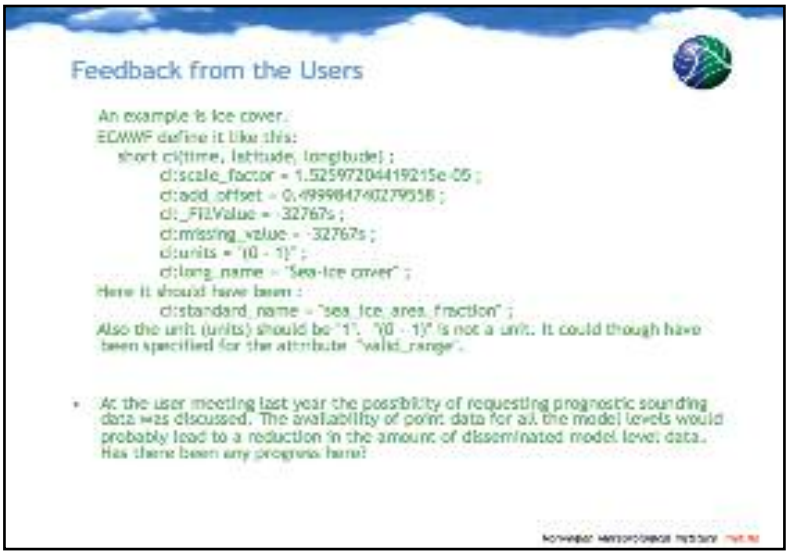

R. Rudsar asked that future ECMWF requests for information allow Computing Representatives more time to gather the information from their users.

ROMANIA

ROMANIA

Roland Cotariu – National Meteorological Administration

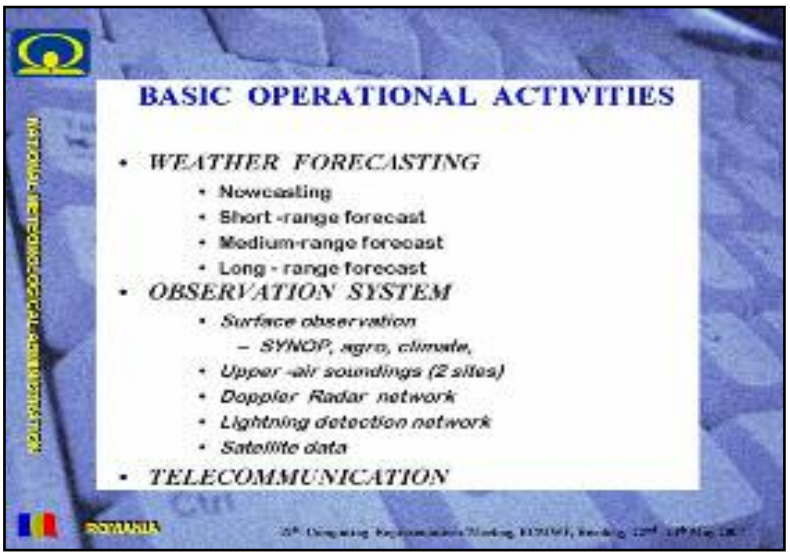

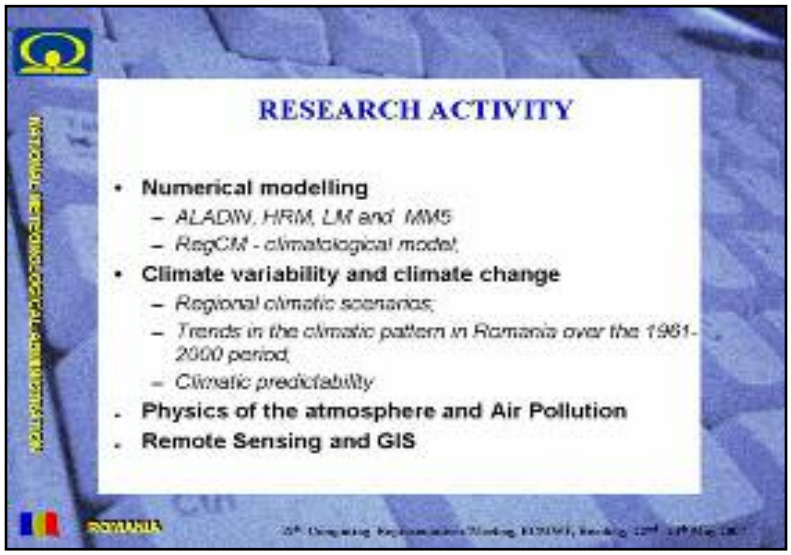

ROMANIA

ROMANIA

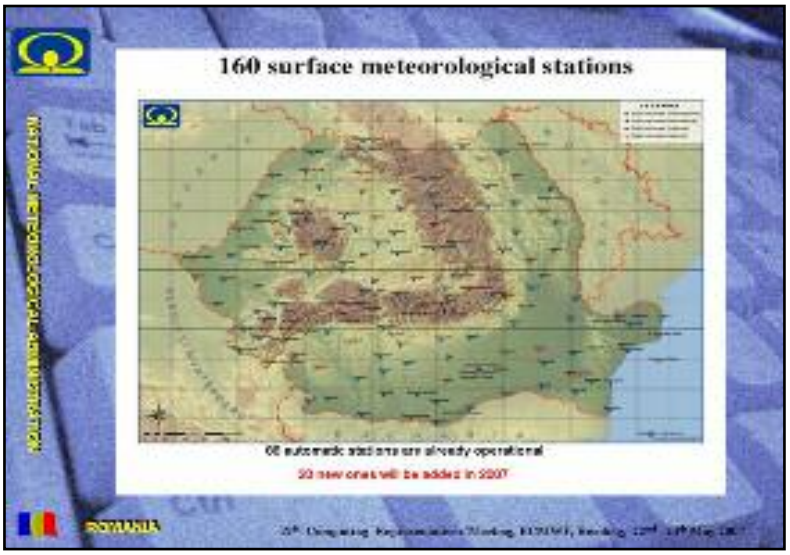

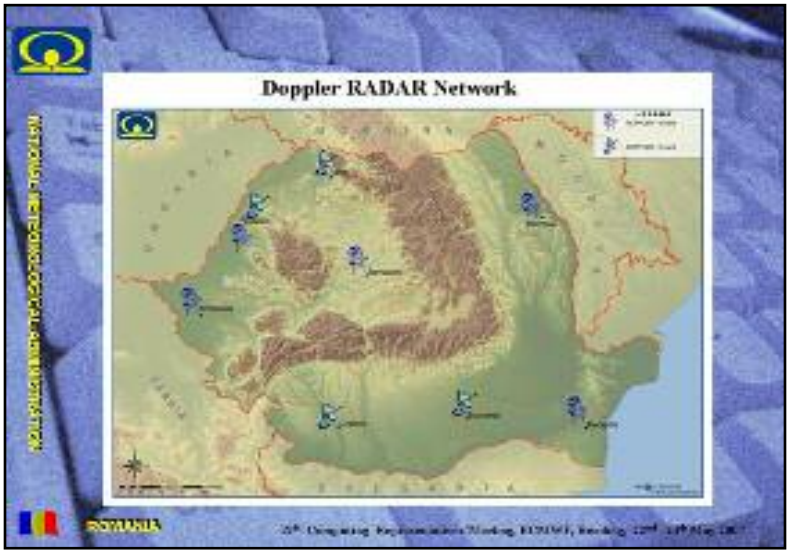

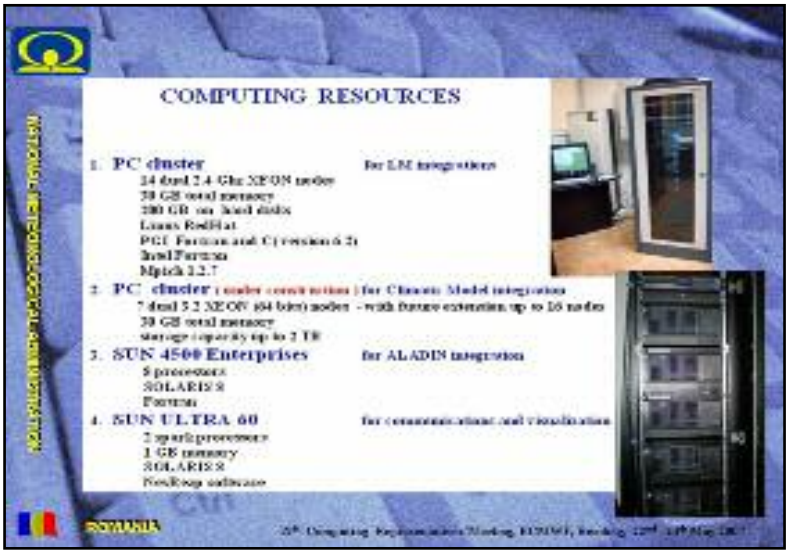

ROMANIA

ROMANIA

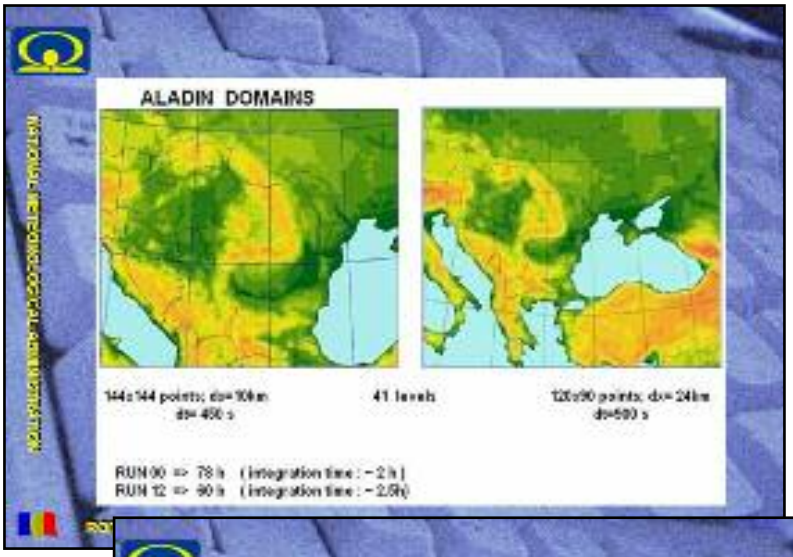

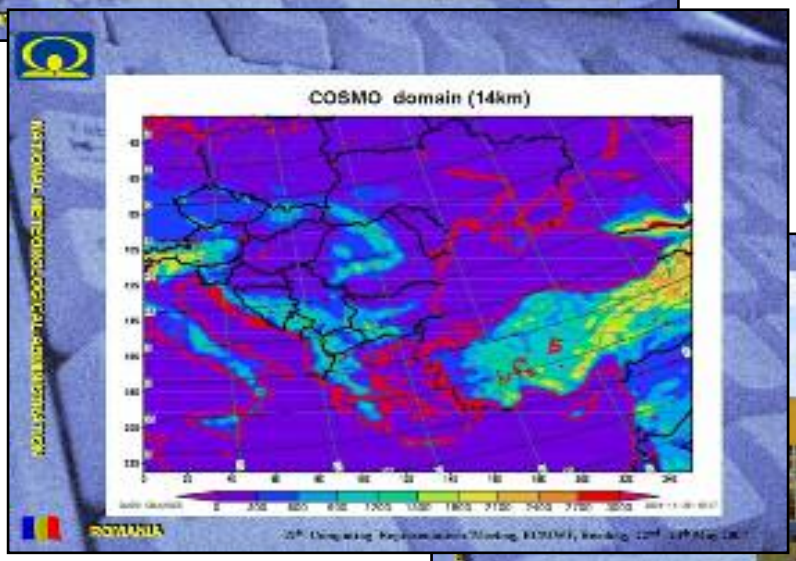

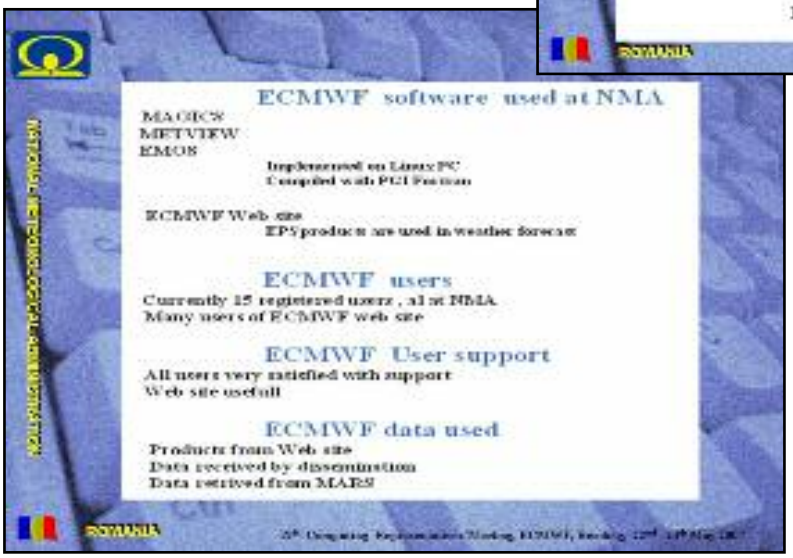

SERBIA SERBIA

Vladimir M. Dimitrijevic – Republic Hydro-Meteorological Service of Serbia

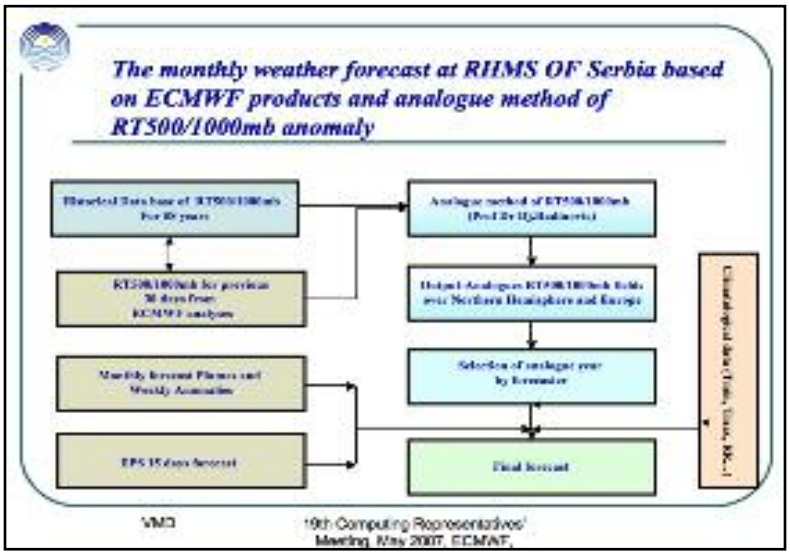

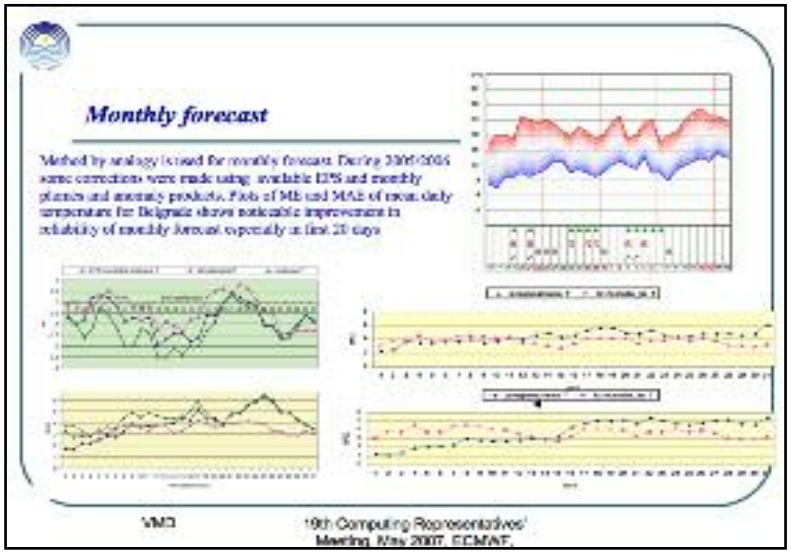

SERBIA

SERBIA

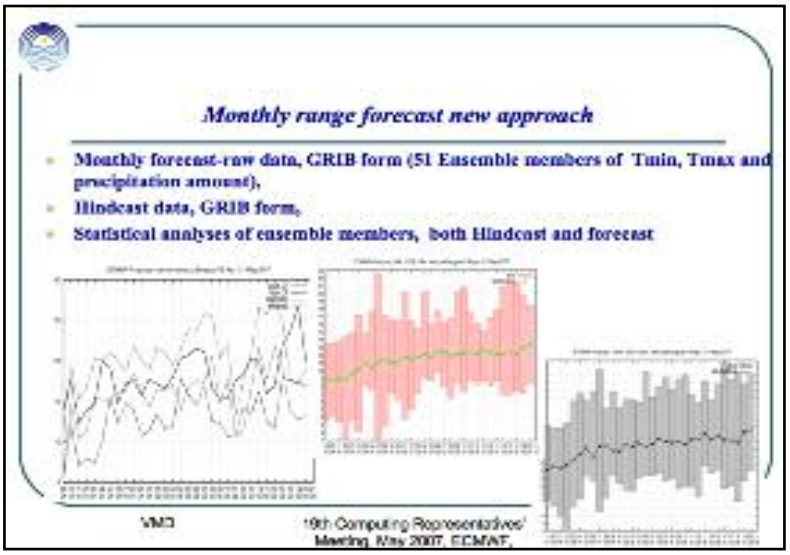

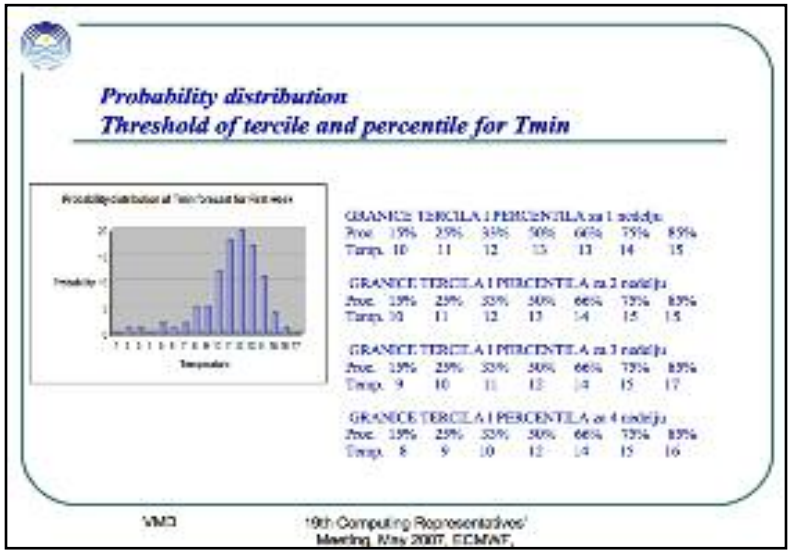

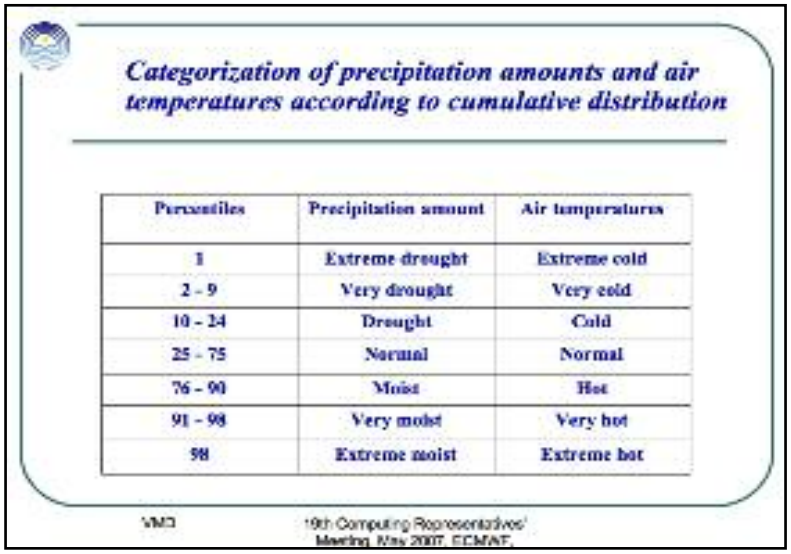

SERBIA

SERBIA

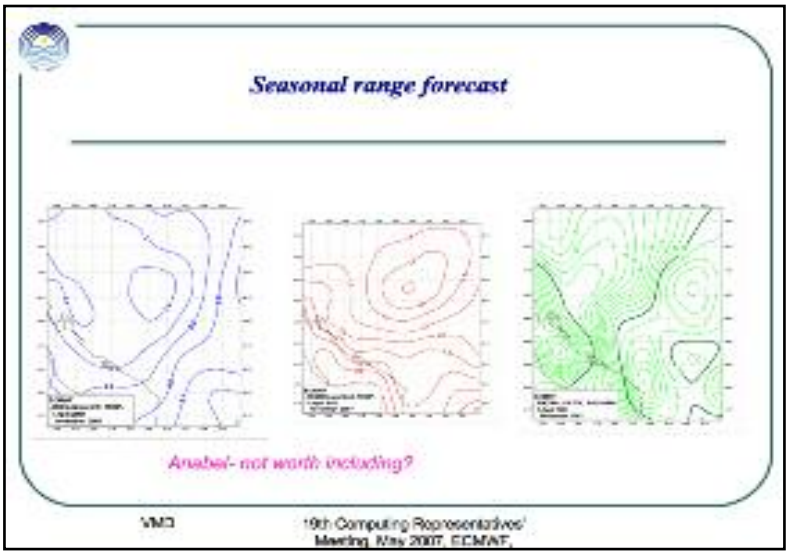

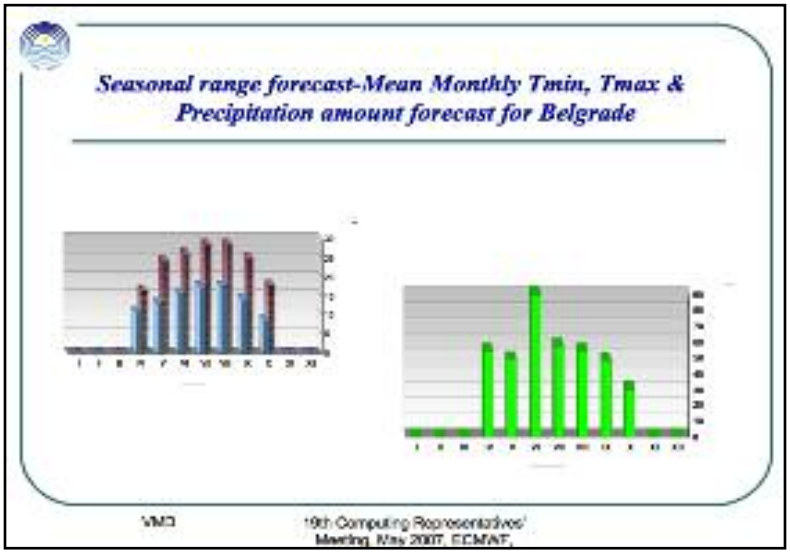

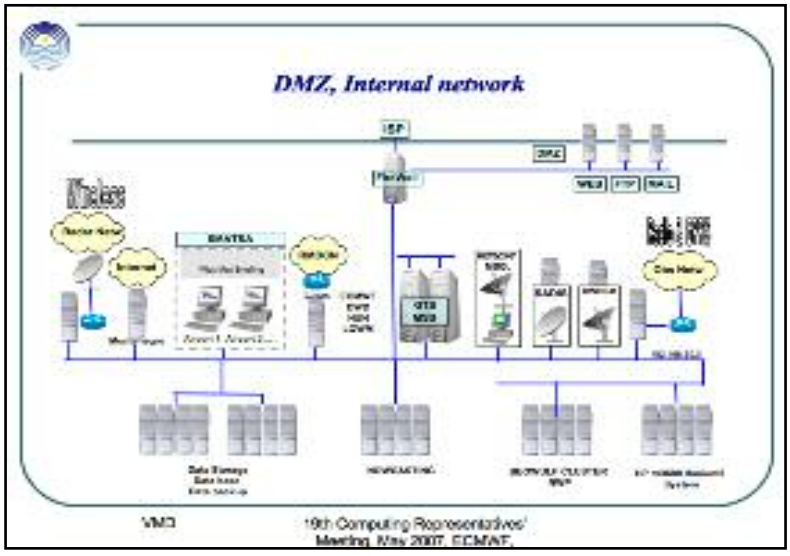

SERBIA

SERBIA

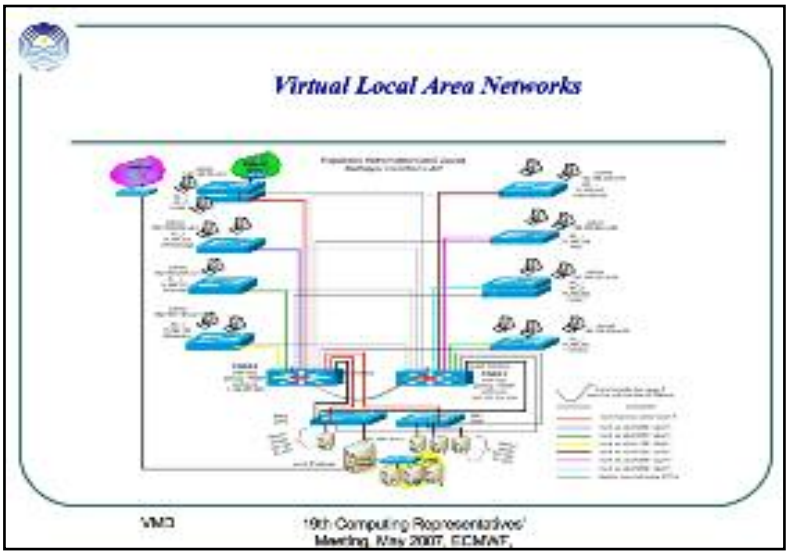

SERBIA

SERBIA

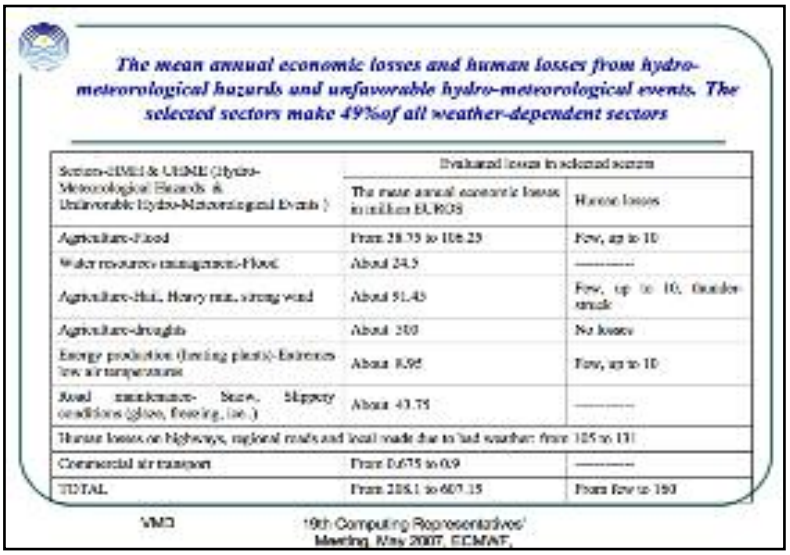

The mean annual economic losses and human losses from hydro-meteorological hazards and unfavorable hydro-meteorological events. The selected sectors make 49%of all weather-dependent sectors

System-HMI & UHME (Hydro-Meteorological Hazards & Unfavorable Hydro-Meteorological Events)	Evaluated losses in selected sectors	
	The mean annual economic losses in million EUROS	Human losses
Agriculture-Flood	From 28.75 to 108.25	Few, up to 10
Water resources management-Flood	About 24.5	----------
Agriculture-Hail, Heavy rain, strong wind	About 51.45	Few, up to 10, Outdoor people
Agriculture-drought	About 300	No losses
Energy production (forcing plants)-Extremes low air temperatures	About 8.95	Few, up to 10
Road maintenance- Snow, Slippery conditions (glaze, freezing, ice..)	About 43.75	----------
Human losses on highways, regional roads and local roads due to bad weather: from 105 to 131		
Commercial air transport	From 0.675 to 0.9	----------
TOTAL	From 228.1 to 607.15	From few to 190

VMO '19th Computing Representatives' Meeting, May 2007, ECMWF,

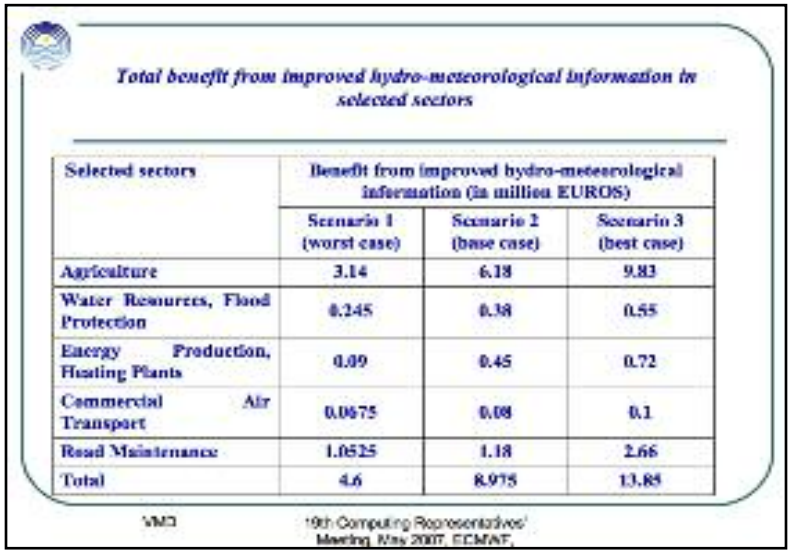

Total benefit from improved hydro-meteorological information in selected sectors

Selected sectors	Benefit from improved hydro-meteorological information (in million EUROS)		
	Scenario 1 (worst case)	Scenario 2 (base case)	Scenario 3 (best case)
Agriculture	3.14	6.18	9.83
Water Resources, Flood Protection	0.245	0.38	0.55
Energy Production, Heating Plants	0.09	0.45	0.72
Commercial Air Transport	0.0675	0.08	0.1
Road Maintenance	1.0525	1.18	2.66
Total	4.6	8.975	13.85

VMO '19th Computing Representatives' Meeting, May 2007, ECMWF,

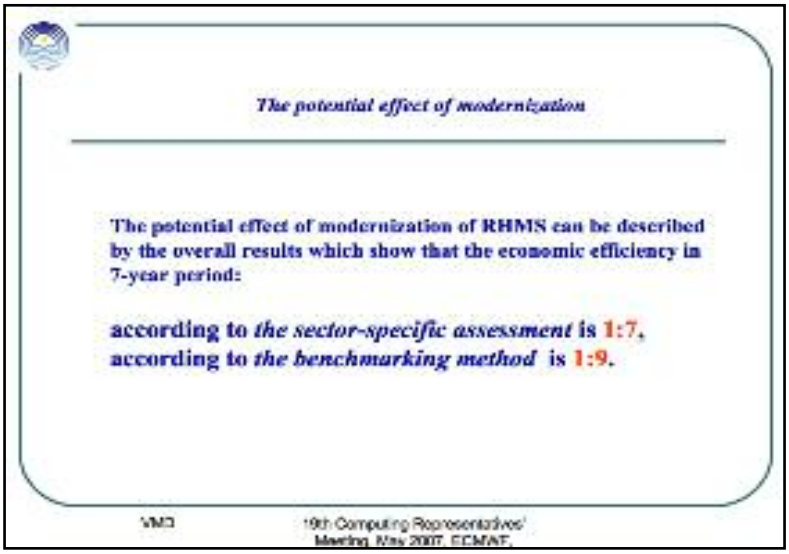

The potential effect of modernization

The potential effect of modernization of RHMS can be described by the overall results which show that the economic efficiency in 7-year period:

according to *the sector-specific assessment* is 1:7,
according to *the benchmarking method* is 1:9.

VMO '19th Computing Representatives' Meeting, May 2007, ECMWF,

SERBIA

SERBIA

SLOVENIA SLOVENIA

Petar Hitij – Ministry for Environment and Spatial Planning, Environmental Agency of the Republic of Slovenia

Virtualization

Virtualization was the main tool for simplifying servers:

- Linux-VServer – more than 15 vservers on 5 hosts, including web server dns&mail, Idap, home directories, nagios ...
- Vm Ware – one host 2 virual servers, including Oracle application server.
- Disk based backup using standard tools like rsync, duplicity.

Linux-VServer

Main features of "Linux-VServer" virtualization:

- single kernel
- excellent (best) performance
- lightweight – best scalability
- simplicity
- easy to install – part of Debian GNU/Linux distribution
- excellent compatibility – works where Linux kernel works
- filesystem sharing without nfs – "mount -bind"
- free

http://www.linux-vserver.org

Processing, Disk and backup capabilities

For this year we plan to aquire:

- new cluster – probably with the newest Intel multicore processors. This time we want to build one ourselves.
- Fibre Channel disk array (SAN). Partially as a disk store for the new cluster and to consolidate disk space across main servers.
- two dedicated backup servers with 5TB SATA disk arrays are already purchased, but not yet operational

R. Urrutia asked why they were using virtualisation for the LINUX machines. *P. Hitij replied that the main reason was security but it also simplified server management.* V. Gislason asked whether they had considered using Xen free virtualisation software instead of Vserver. *P. Hitij replied that it had been considered but Vserver was simpler to implement; he also believed that Xen used more than one kernel*

SPAIN SPAIN

Eduardo Monreal – Instituto Naçional de Meteorología, Spain

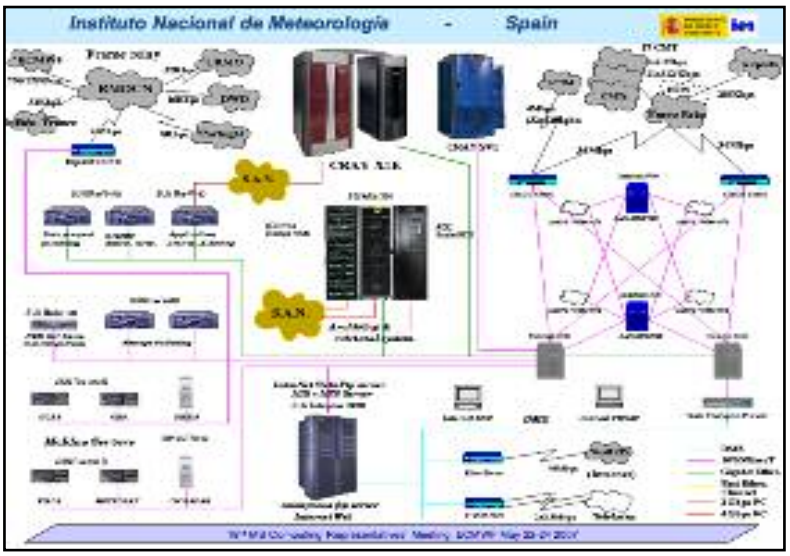

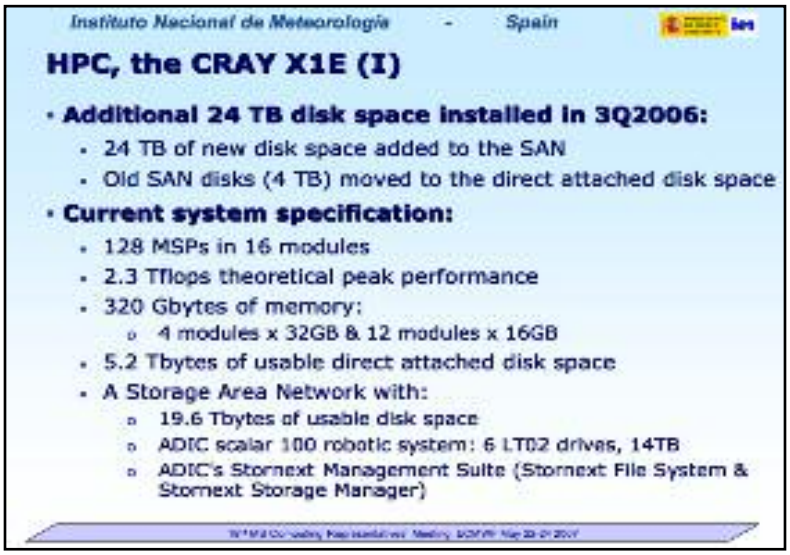

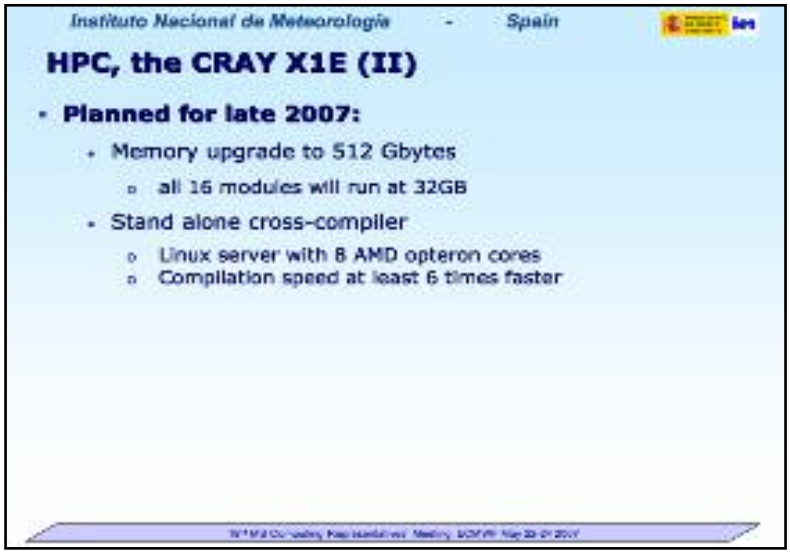

SPAIN

SPAIN

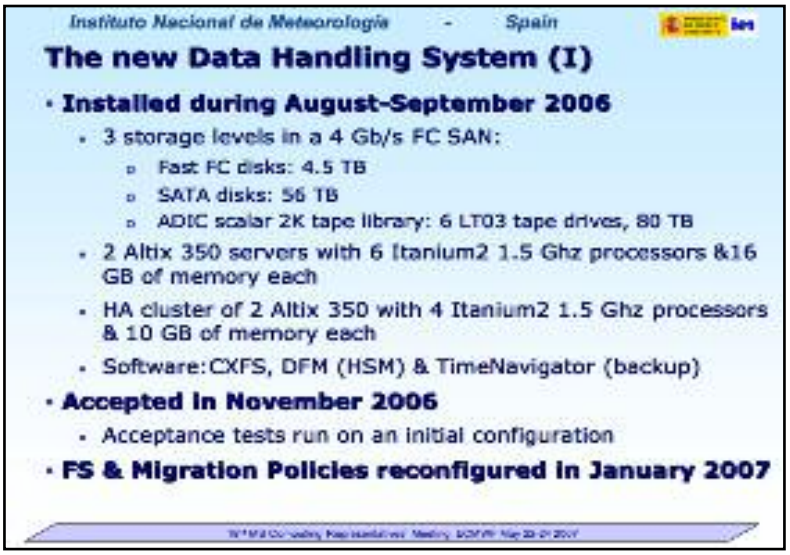
Instituto Nacional de Meteorologia - Spain
The new Data Handling System (I)
• Installed during August-September 2006
• 3 storage levels in a 4 Gb/s FC SAN:
Fast FC disks: 4.5 TB
SATA disks: 56 TB
ADIC scalar 2K tape library: 6 LTO3 tape drives, 80 TB
• 2 Altix 350 servers with 6 Itanium2 1.5 Ghz processors &16 GB of memory each
• HA cluster of 2 Altix 350 with 4 Itanium2 1.5 Ghz processors & 10 GB of memory each
• Software:CXFS, DFM (HSM) & TimeNavigator (backup)
• Accepted in November 2006
• Acceptance tests run on an initial configuration
• FS & Migration Policies reconfigured in January 2007

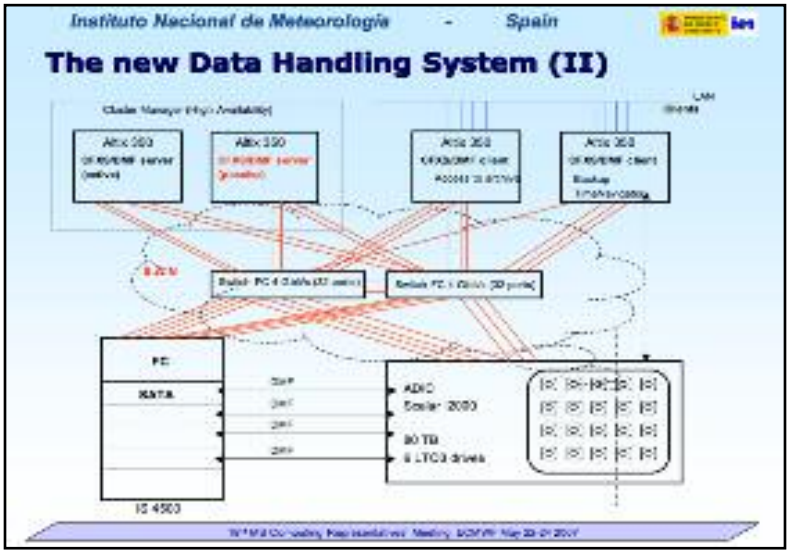
Instituto Nacional de Meteorologia - Spain
The new Data Handling System (II)
FC
SATA
ADIC
Scalar 2000
80 TB
6 LTO3 drives
IS 4500

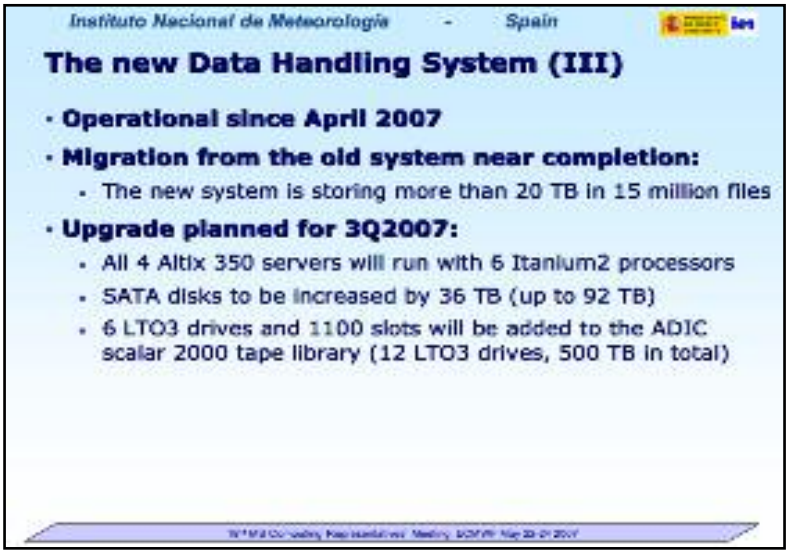
Instituto Nacional de Meteorologia - Spain
The new Data Handling System (III)
• Operational since April 2007
• Migration from the old system near completion:
• The new system is storing more than 20 TB in 15 million files
• Upgrade planned for 3Q2007:
• All 4 Altix 350 servers will run with 6 Itanium2 processors
• SATA disks to be increased by 36 TB (up to 92 TB)
• 6 LTO3 drives and 1100 slots will be added to the ADIC scalar 2000 tape library (12 LTO3 drives, 500 TB in total)

SPAIN

SPAIN

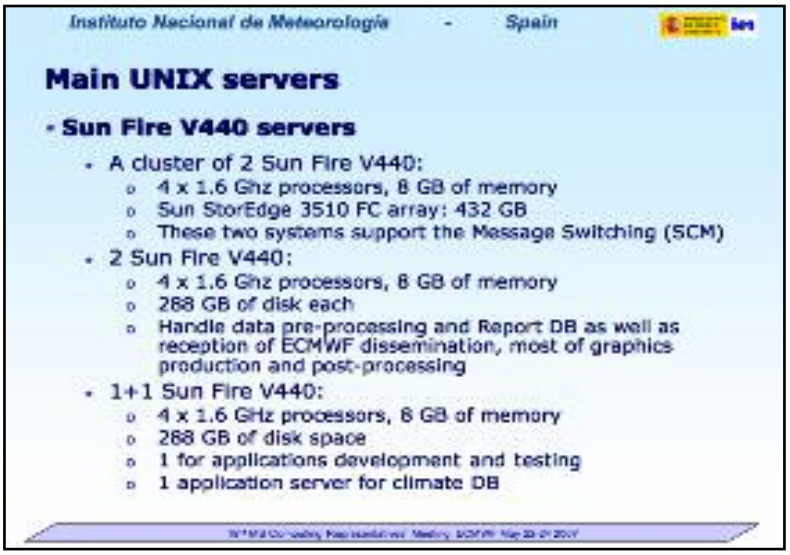
Instituto Nacional de Meteorologia - Spain
Main UNIX servers
Sun Fire V440 servers
A cluster of 2 Sun Fire V440:
4 x 1.6 Ghz processors, 8 GB of memory
Sun StorEdge 3510 FC array: 432 GB
These two systems support the Message Switching (SCM)
2 Sun Fire V440:
4 x 1.6 Ghz processors, 8 GB of memory
288 GB of disk each
Handle data pre-processing and Report DB as well as reception of ECMWF dissemination, most of graphics production and post-processing
1+1 Sun Fire V440:
4 x 1.6 GHz processors, 8 GB of memory
288 GB of disk space
1 for applications development and testing
1 application server for climate DB

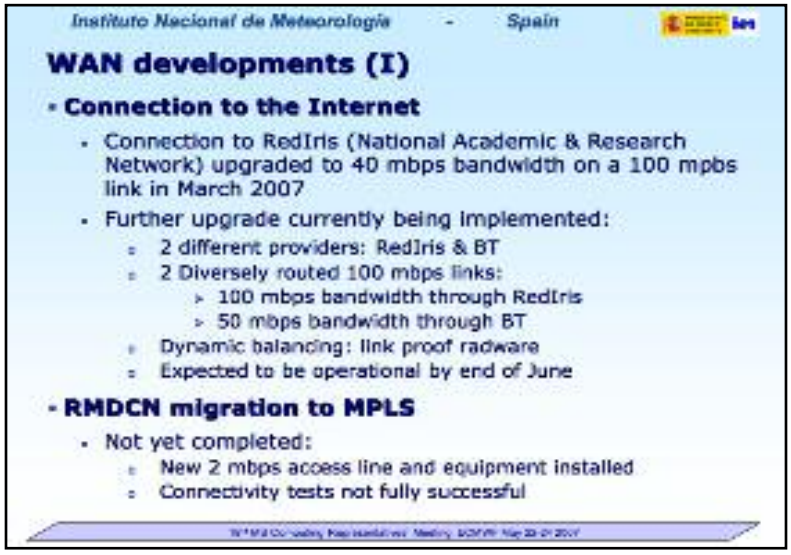
Instituto Nacional de Meteorologia - Spain
WAN developments (I)
Connection to the Internet
Connection to RedIris (National Academic & Research Network) upgraded to 40 mbps bandwidth on a 100 mpbs link in March 2007
Further upgrade currently being implemented:
2 different providers: RedIris & BT
2 Diversely routed 100 mbps links:
100 mbps bandwidth through RedIris
50 mbps bandwidth through BT
Dynamic balancing: link proof radware
Expected to be operational by end of June
RMDCN migration to MPLS
Not yet completed:
New 2 mbps access line and equipment installed
Connectivity tests not fully successful

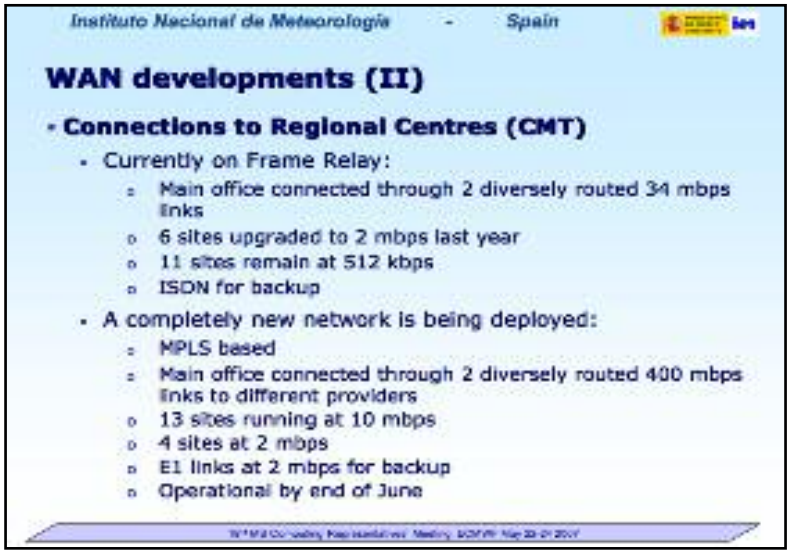
Instituto Nacional de Meteorologia - Spain
WAN developments (II)
Connections to Regional Centres (CMT)
Currently on Frame Relay:
Main office connected through 2 diversely routed 34 mbps links
6 sites upgraded to 2 mbps last year
11 sites remain at 512 kbps
ISDN for backup
A completely new network is being deployed:
MPLS based
Main office connected through 2 diversely routed 400 mbps links to different providers
13 sites running at 10 mbps
4 sites at 2 mbps
E1 links at 2 mbps for backup
Operational by end of June

SPAIN

SPAIN

Instituto Nacional de Meteorologia - Spain

Use of ECMWF systems

- **Registered users:**
  - At present, 77 Spanish users are registered (72 last year)
    - 63 from INM and 14 from universities
    - 6 not active in 2006-2007
    - 59 logged on in April & May
  - 10 more Spanish users in 3 Special Projects
- **Work done is for the most part MARS data retrievals, particularly access to ERA-40 dataset**
- **Metview used in batch mode to produce derived EPS products**
- **All users use the Internet for data transfer**
  - Data transfer over RMDCN only for dissemination and other operational use

Instituto Nacional de Meteorologia - Spain

Use of HPCF

- **Last year 18 users used the HPCF. They basically worked in the following areas:**
  - HIRLAM model runs using the reference system
  - Trajectory computations
  - Studies on Climate variability
  - Statistical downscaling of seasonal forecast outputs
  - Integration of RCA/HIRLAM within EU ENSEMBLES Project framework
- **HIRLAM accounts for the 96% of 2006 used allocation:**
  - Data assimilation group (4D-Var): 46%
  - Parallel runs: 36%
  - Model development: 14%
- **A 2% due to Seasonal FC downscaling**

Instituto Nacional de Meteorologia - Spain

HPCF: CPU usage

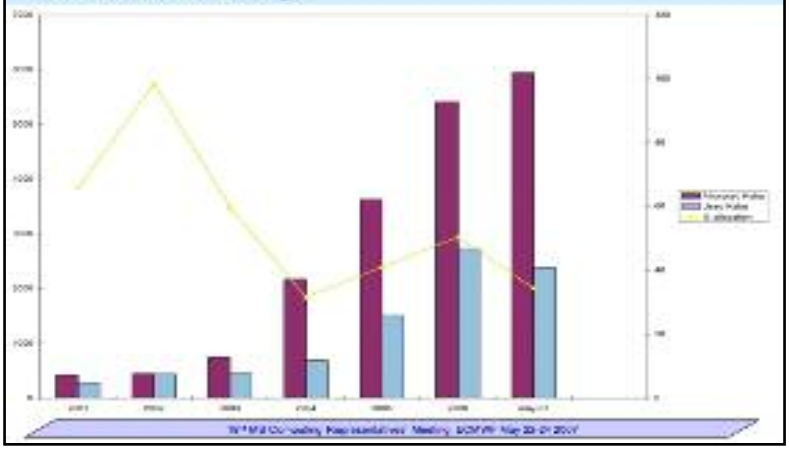

SPAIN

SPAIN

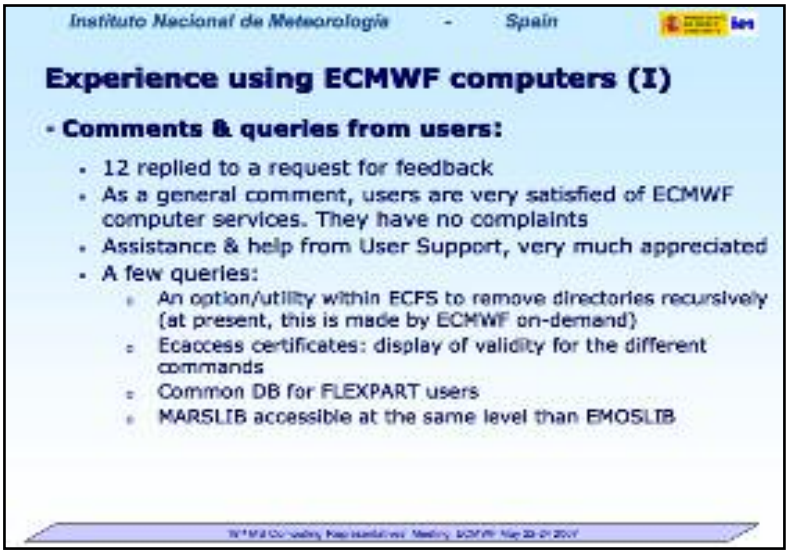
Instituto Nacional de Meteorología - Spain
Experience using ECMWF computers (I)
Comments & queries from users:
12 replied to a request for feedback
As a general comment, users are very satisfied of ECMWF computer services. They have no complaints
Assistance & help from User Support, very much appreciated
A few queries:
An option/utility within ECFS to remove directories recursively (at present, this is made by ECMWF on-demand)
Ecaccess certificates: display of validity for the different commands
Common DB for FLEXPART users
MARSLIB accessible at the same level than EMOSLIB

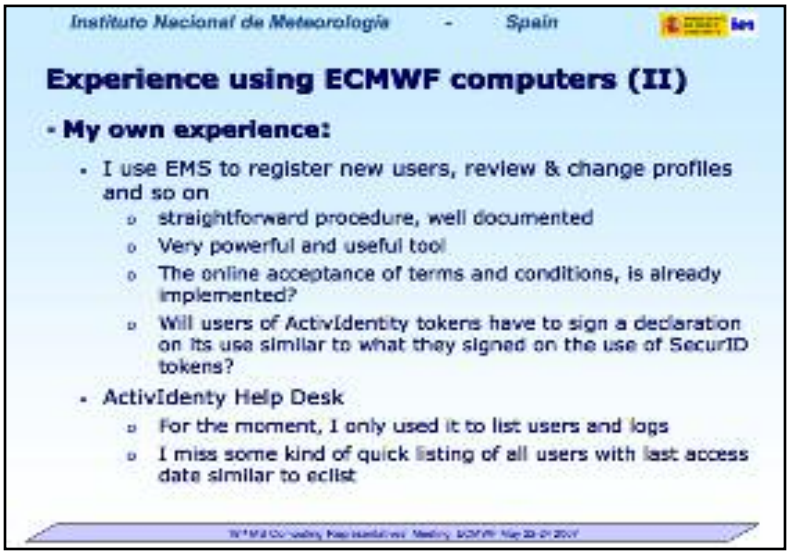
Instituto Nacional de Meteorología - Spain
Experience using ECMWF computers (II)
My own experience:
I use EMS to register new users, review & change profiles and so on
straightforward procedure, well documented
Very powerful and useful tool
The online acceptance of terms and conditions, is already implemented?
Will users of ActivIdentity tokens have to sign a declaration on its use similar to what they signed on the use of SecurID tokens?
ActivIdenty Help Desk
For the moment, I only used it to list users and logs
I miss some kind of quick listing of all users with last access date similar to eclist

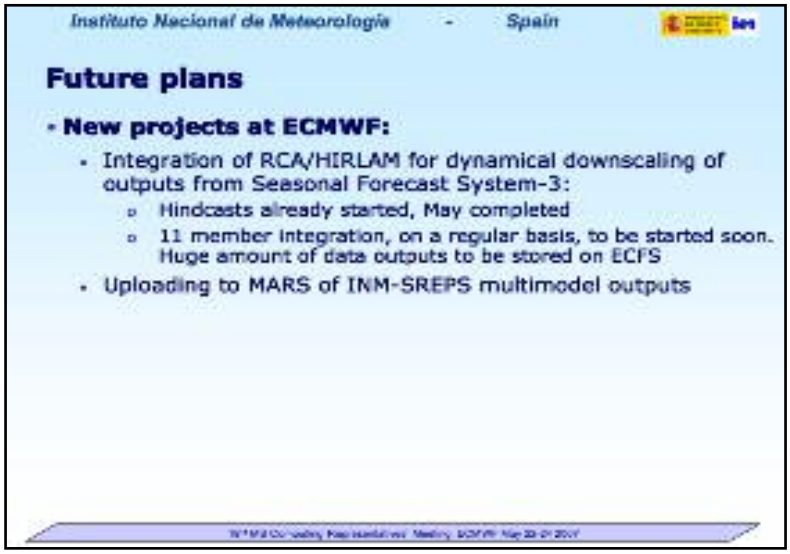
Instituto Nacional de Meteorología - Spain
Future plans
New projects at ECMWF:
Integration of RCA/HIRLAM for dynamical downscaling of outputs from Seasonal Forecast System-3:
Hindcasts already started, May completed
11 member integration, on a regular basis, to be started soon. Huge amount of data outputs to be stored on ECFS
Uploading to MARS of INM-SREPS multimodel outputs

SPAIN SPAIN

In response to E. Monreal's query, in relation to EMS, about on line acceptance of ECMWF's Terms and Conditions of use, U. Modigliani stated that, as a first step, every user gets a link to a web page containing all the Terms and Conditions as part of the registration process. Currently, any new users of ActivIdentity cards need to sign a declaration similar to the one used with the SecurID cards.

D. Lucas noted that MétéoFrance asks all their users (internal and external) to sign a licence agreement with MétéoFrance, so any new procedure to be introduced would need to take account of this. D. Garçon noted that a facility to allow a quick listing of all users with their date of last access should be easy to create.

SWEDEN SWEDEN

Rafael Urrutia – Swedish Meteorological and Hydrological Institute (SMHI)

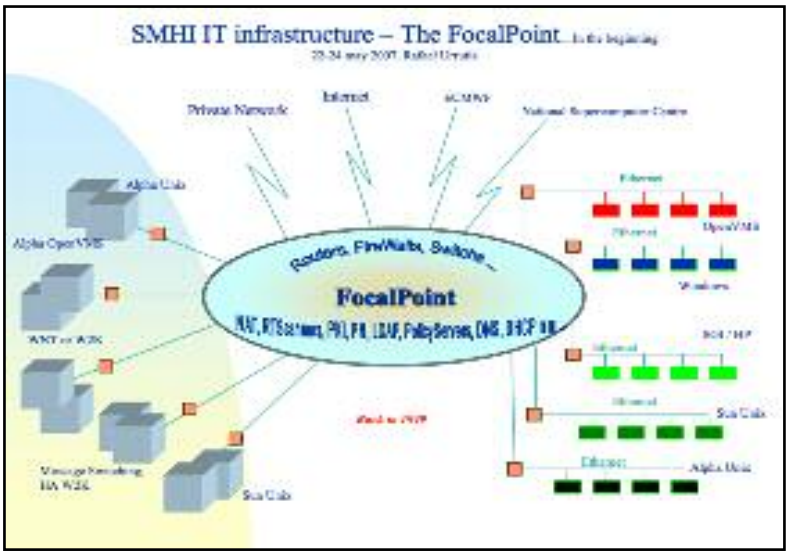

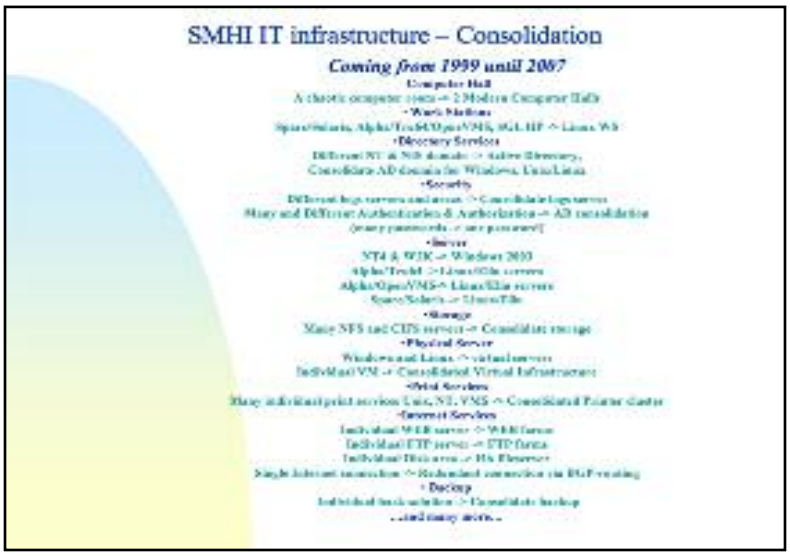

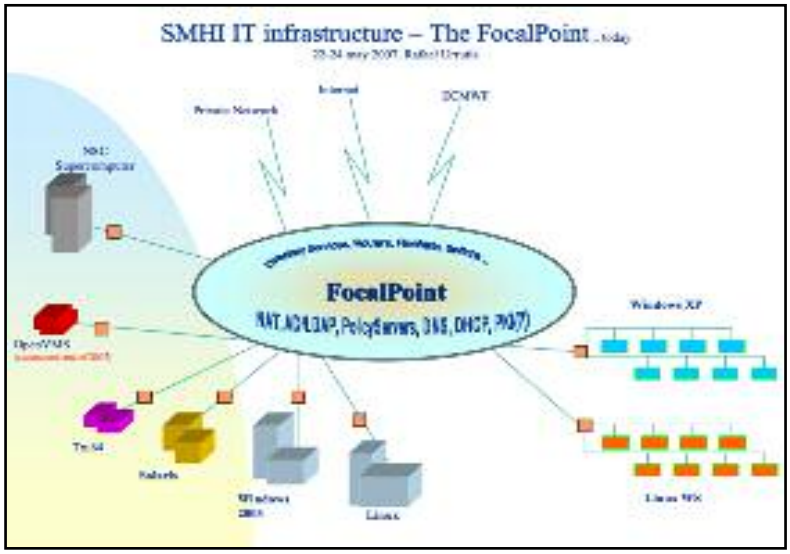

SWEDEN

SWEDEN

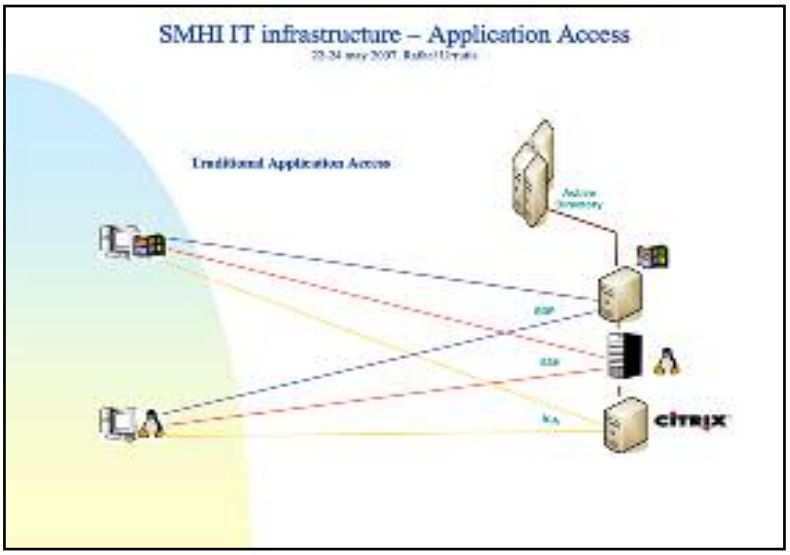

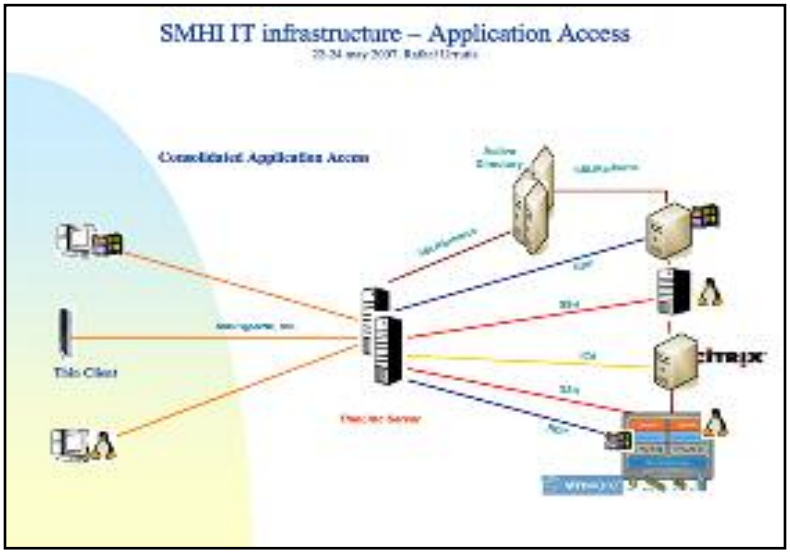

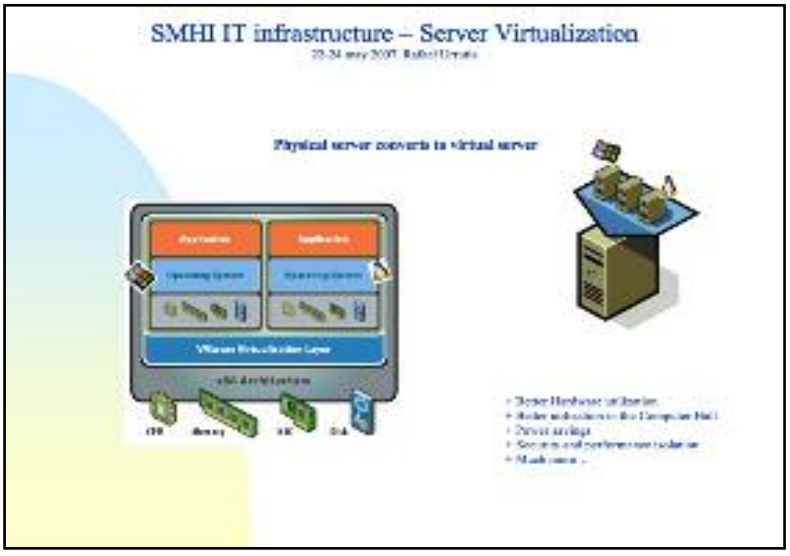

SWEDEN

SWEDEN

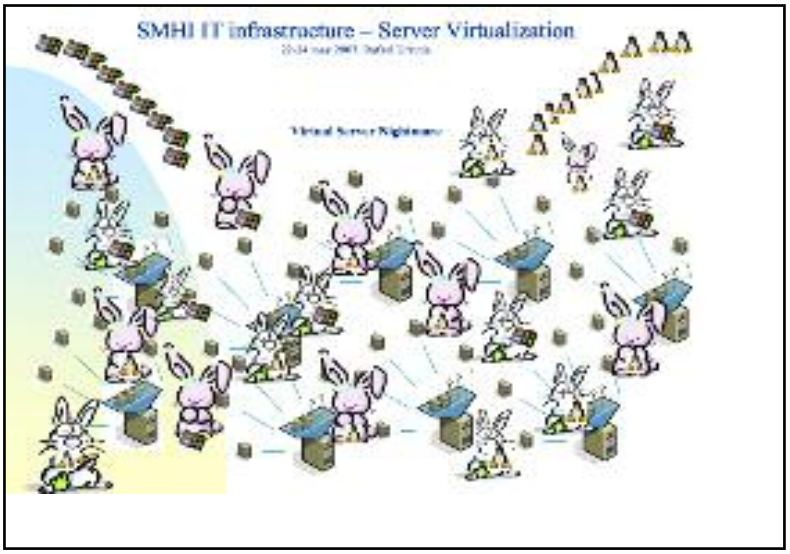

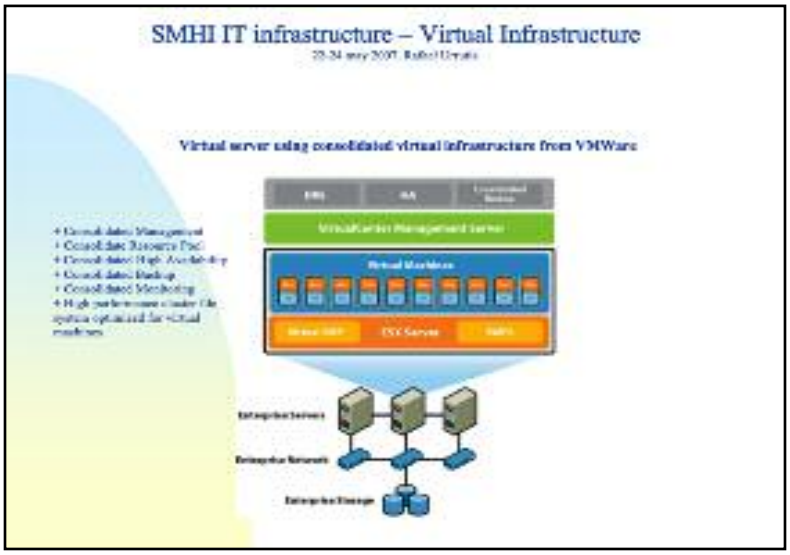

SWITZERLAND SWITZERLAND

Peter Roth – Federal Office of Meteorology and Climatology, MeteoSwiss

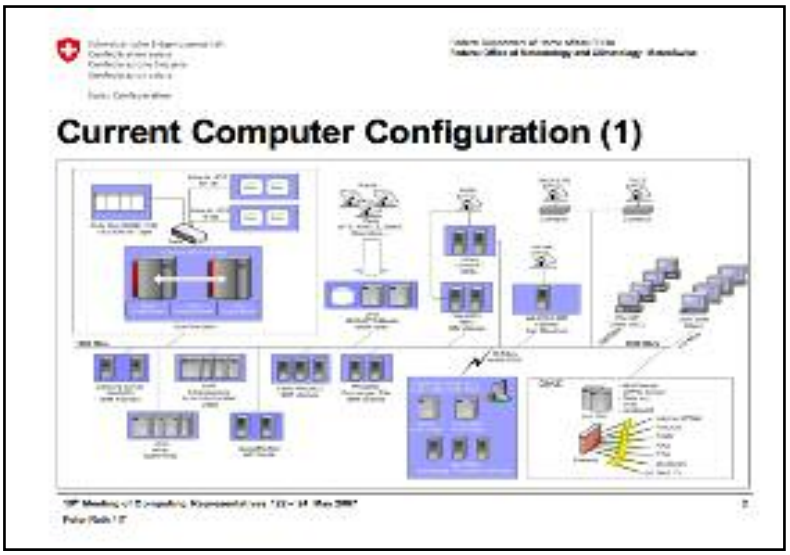

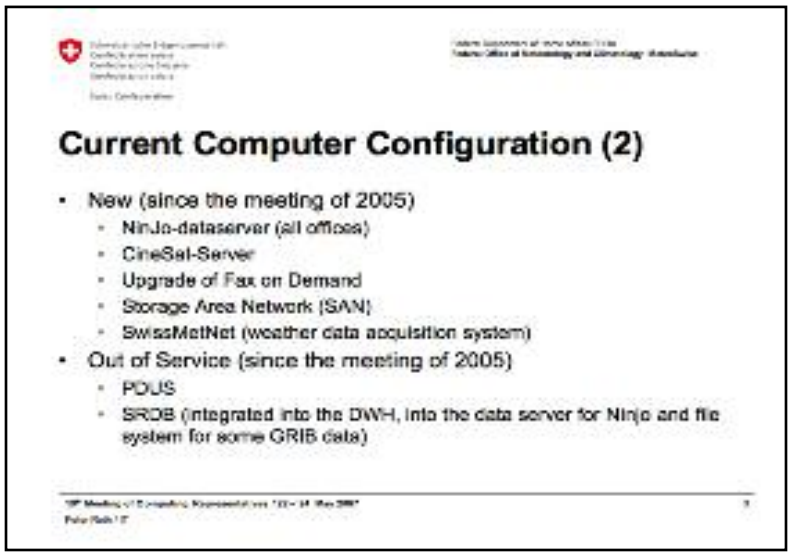

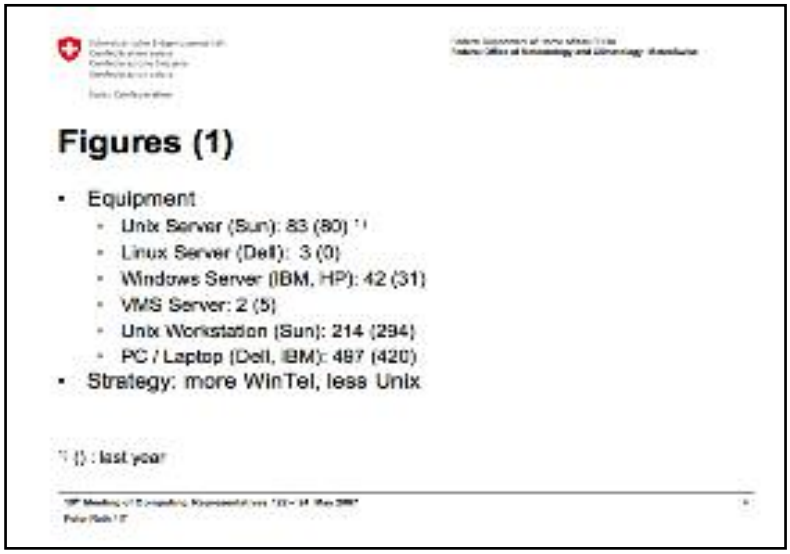

SWITZERLAND

SWITZERLAND

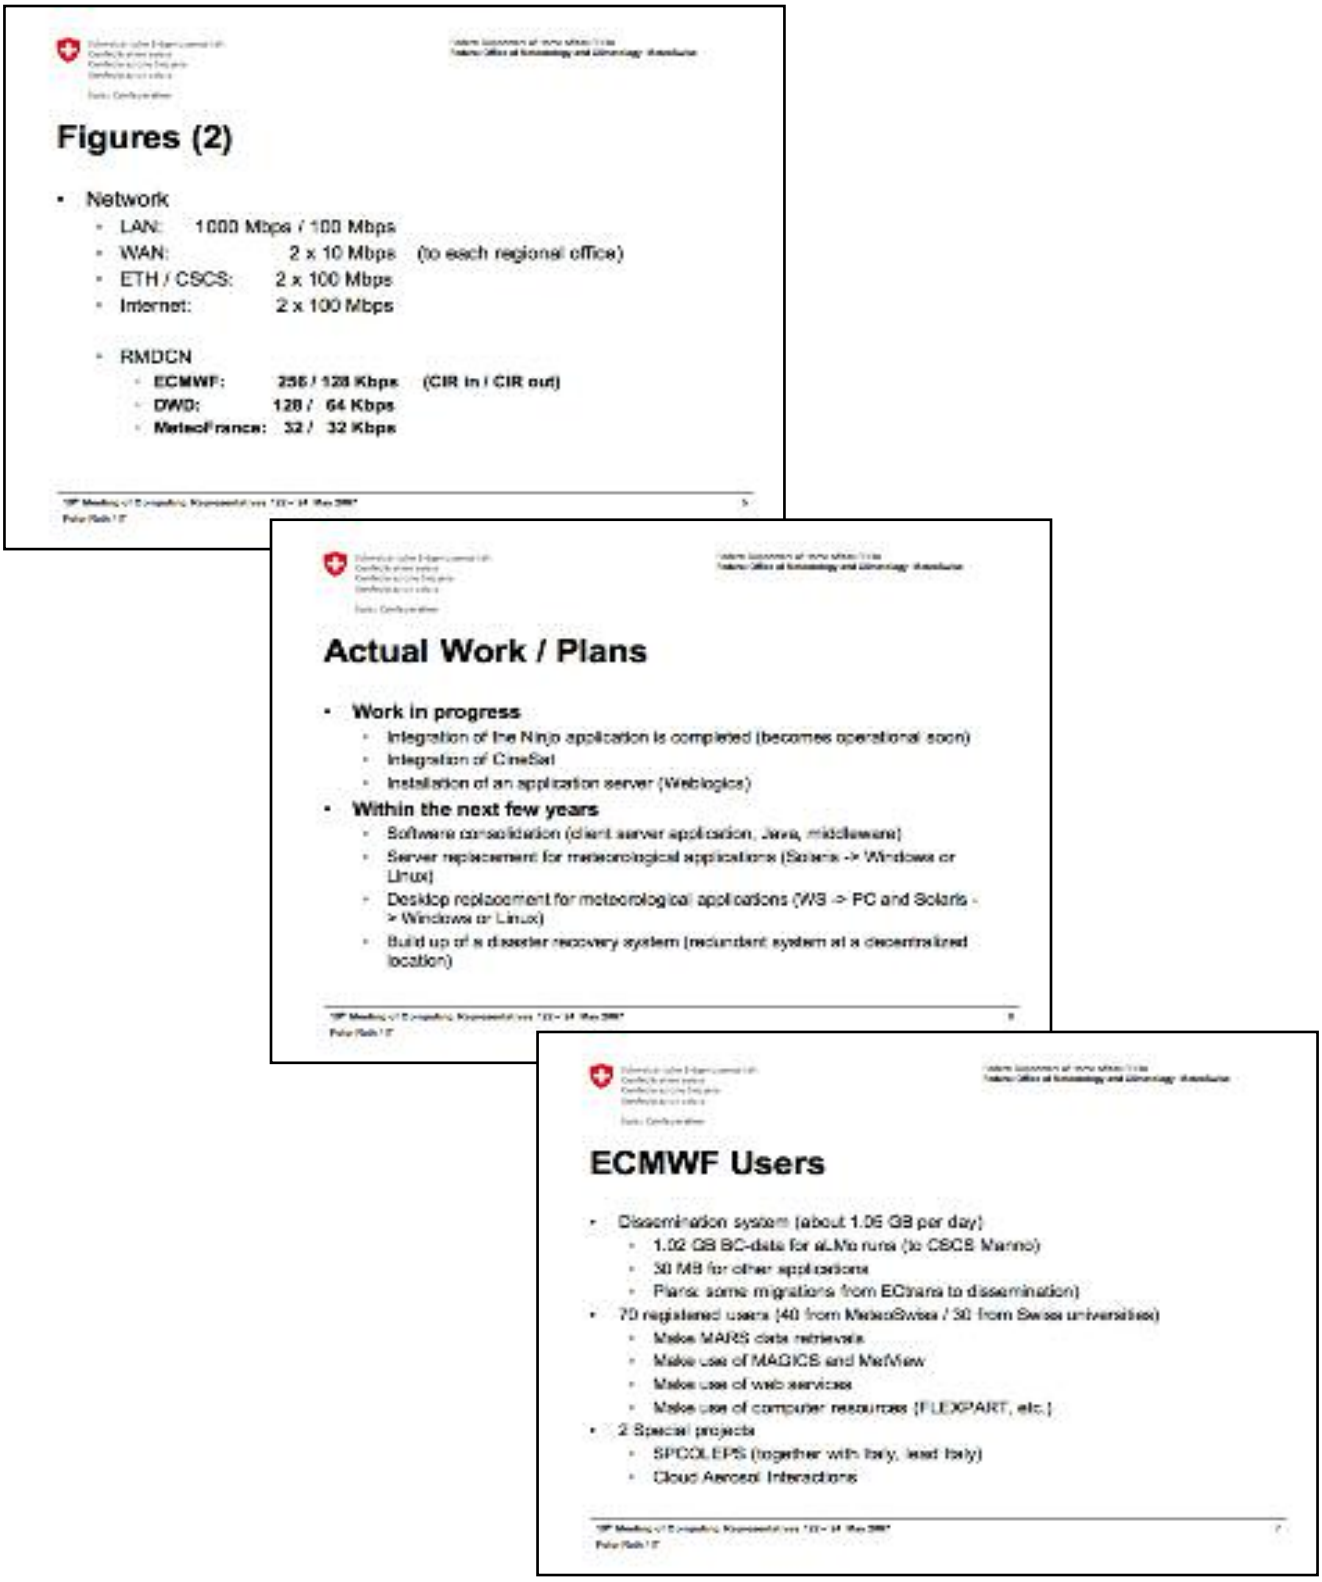

R. Urrutia asked why they used Windows. *P. Roth replied that it enabled them to follow DWD and benefit from their experience.*

UNITED KINGDOM

UNITED KINGDOM

Richard Barnes – Met Office, United Kingdom

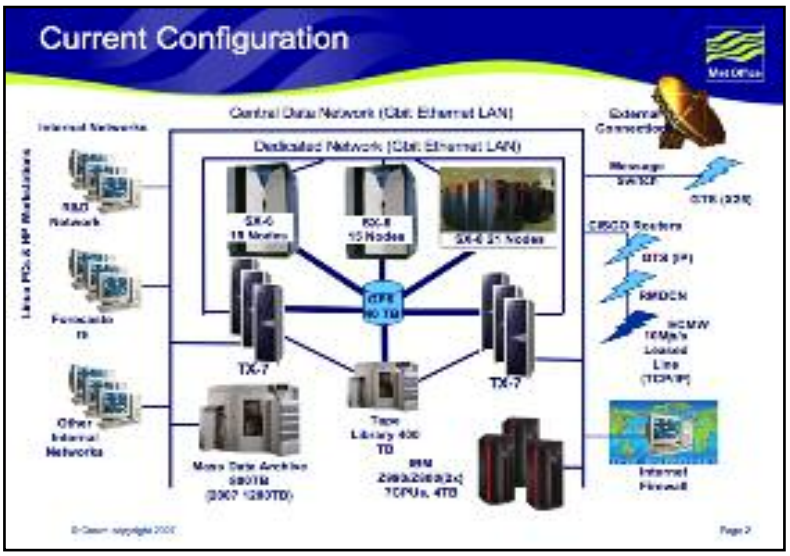

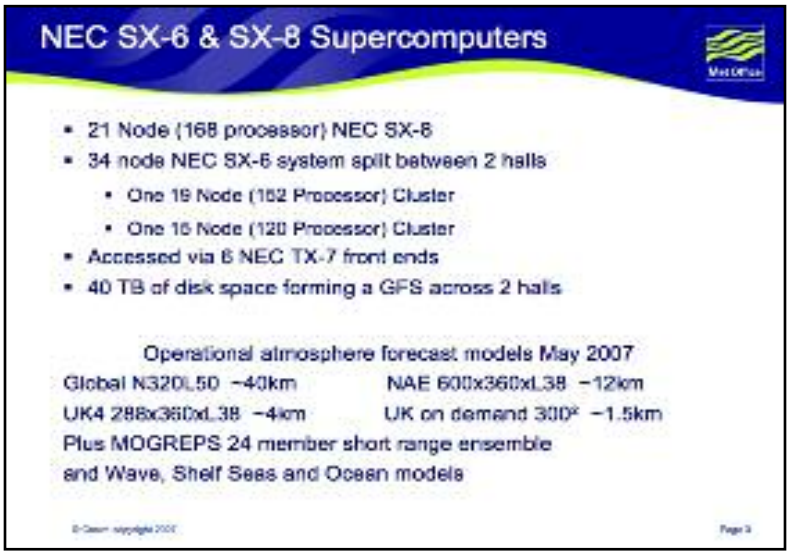

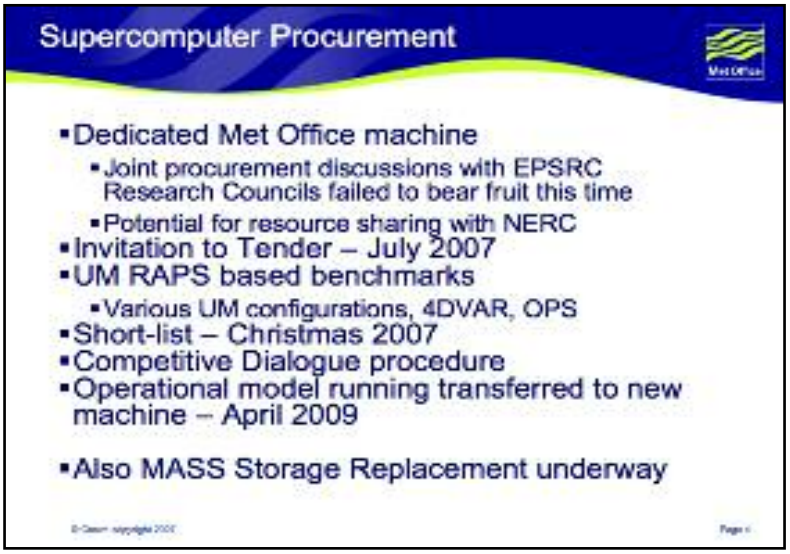

UNITED KINGDOM

UNITED KINGDOM

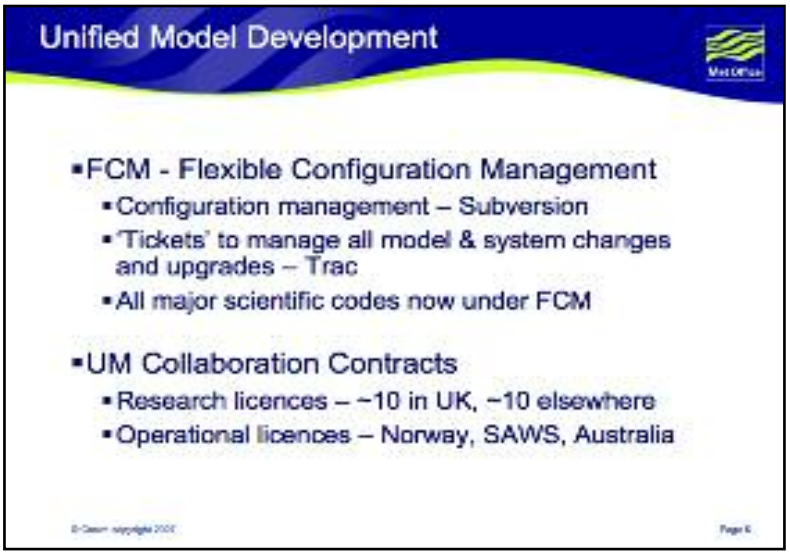

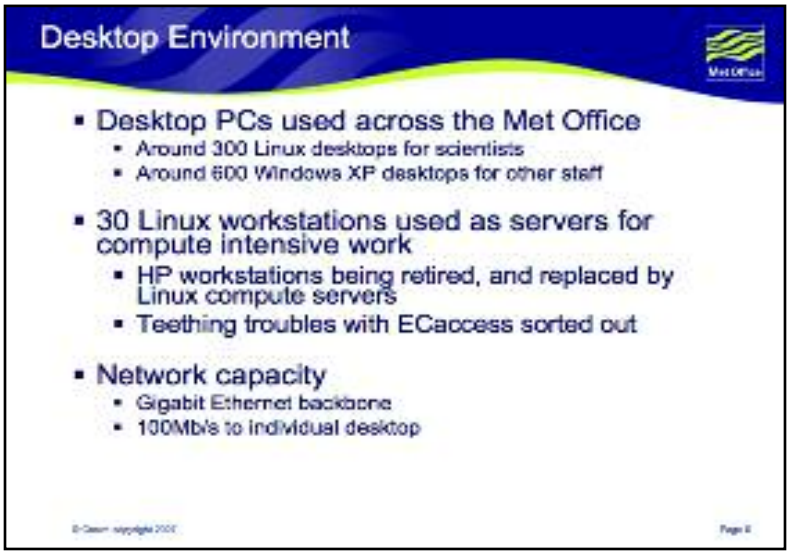

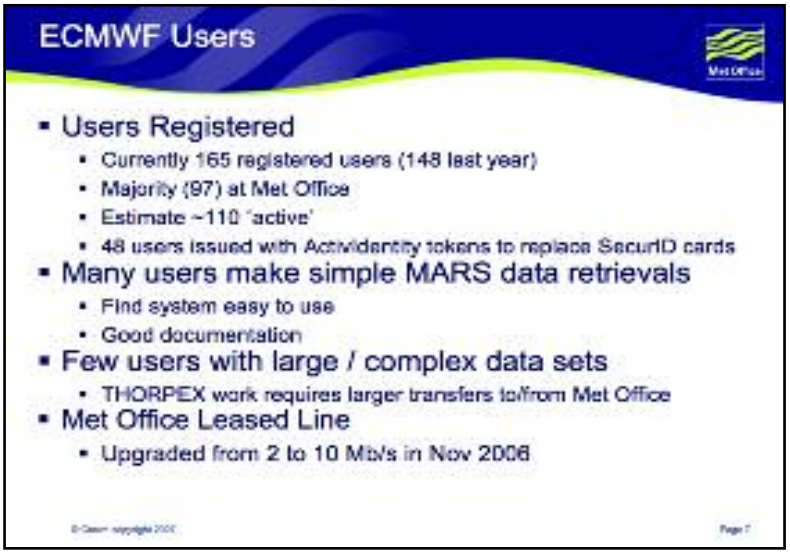

UNITED KINGDOM

UNITED KINGDOM

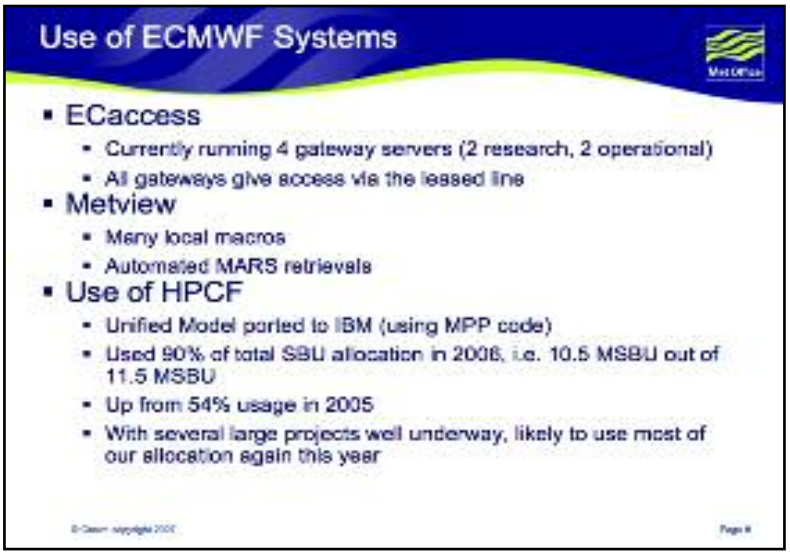

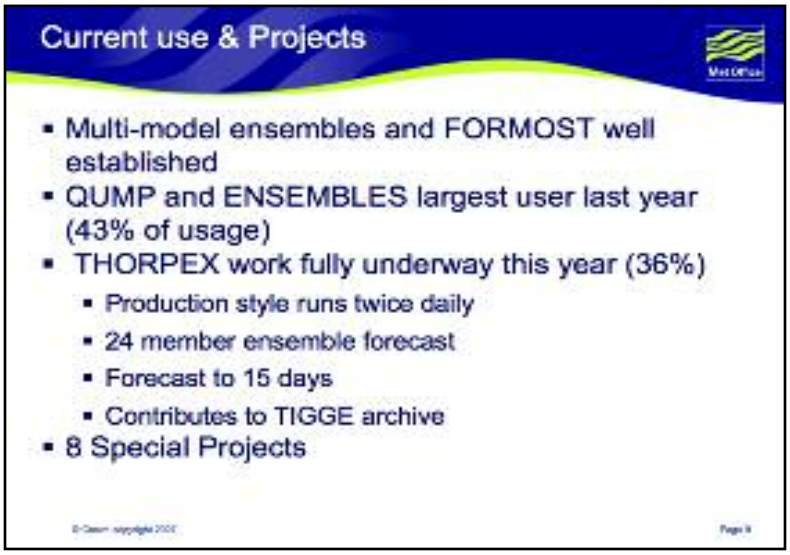

UNITED KINGDOM

UNITED KINGDOM

Met Office THORPEX medium-range ensemble

- Global 15-day ensemble forecasts, based on MOGREPS (Met Office Global and Regional Ensemble Prediction System).
  - Initial conditions are produced at Met Office, then transferred to ECMWF for running the ensemble forecasts.
  - 24 members (control + 23 ETKF-based perturbations), run twice a day (0 and 12 UTC).
  - Resolution: N144 (0.833° x 1.25°), 38 levels
  - The ensemble suite is being further developed to run as a "time critical" suite.
  - Output from the ensemble forecasts is now being added to the TIGGE (THORPEX Interactive Grand Global Ensemble) database hosted by ECMWF, and also copied to NCAR (and, in future, CMA).

Products from Met Office medium-range ensemble

- Output from the ensemble forecasts is post-processed both on ecgate, and at the Met Office, to generate experimental forecast products, particularly to highlight predicted high-impact weather.
- The example below shows the probability of strong winds (>60kt at 925 hPa) for the storm of 18th January 2007, at forecast ranges from 1 day to 14 days.

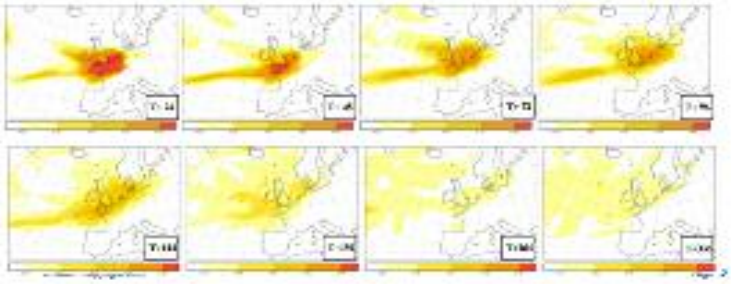

CTBTO

CTBTO

Gerhard Wotawa – Preparatory Commission for the Comprehensive Nuclear-Test-Ban Treaty Organization, Provisional Technical Secretariat, Vienna, Austria

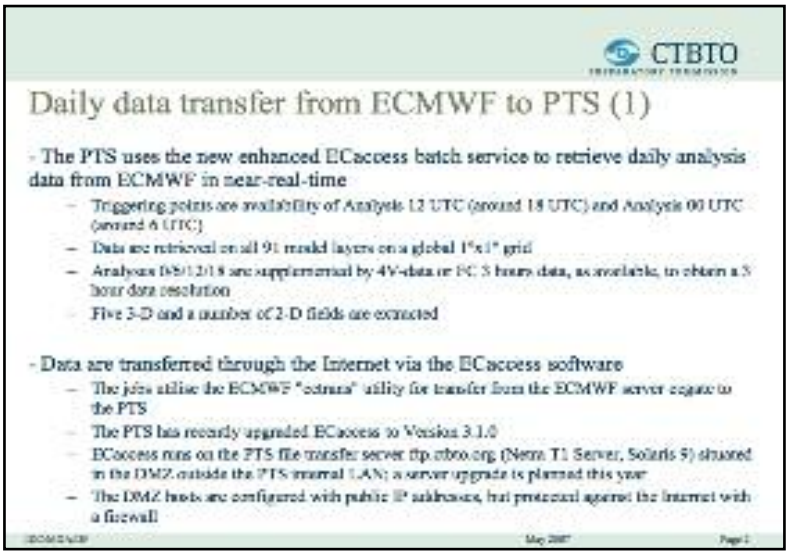

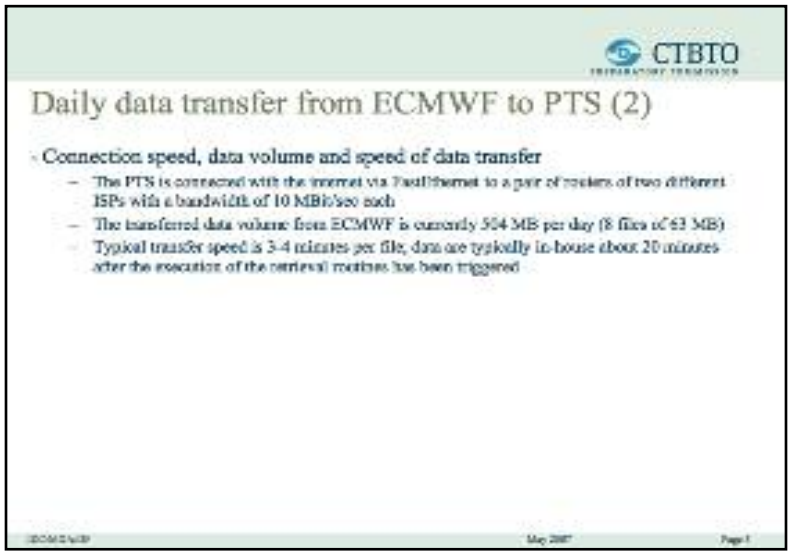

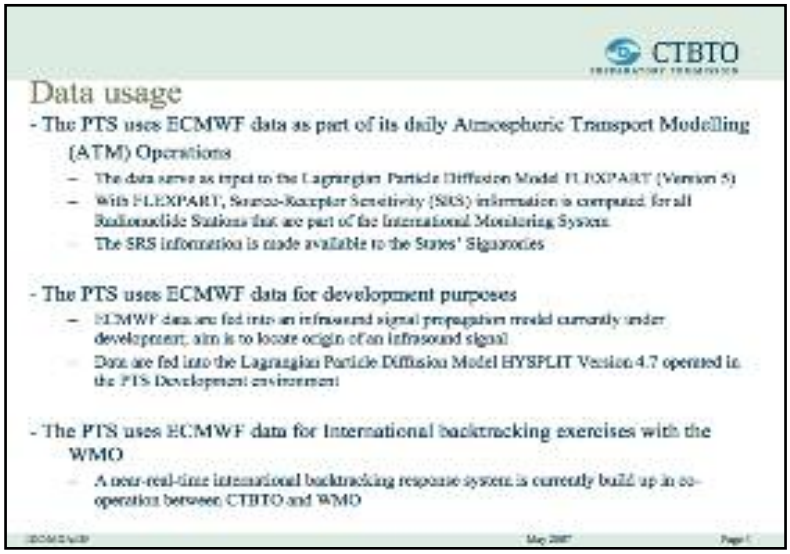

CTBTO

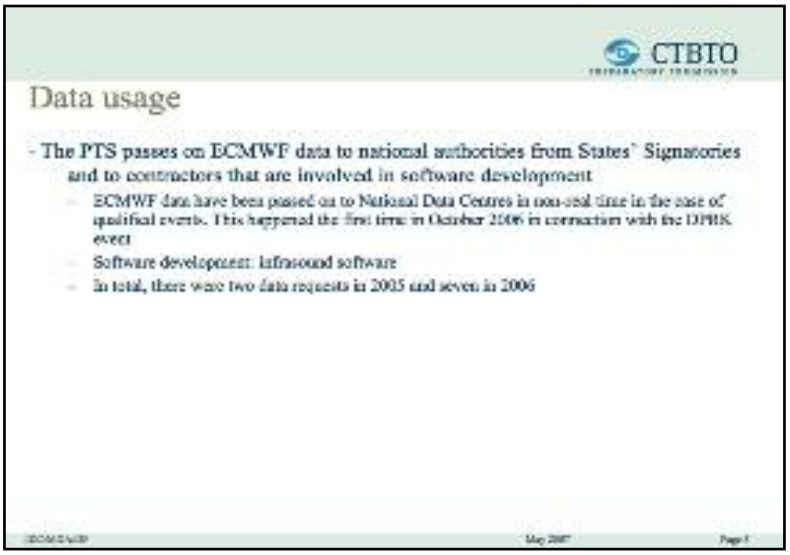

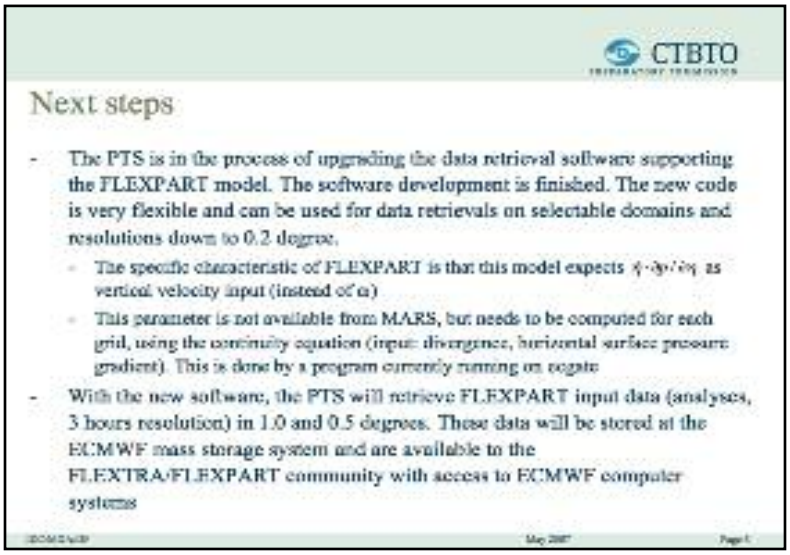

T. Lorenzen assumed that the CTBTO runs mainly retrospective ATM models but asked whether it also receives ECMWF products out to day 15 to study how events are likely to evolve. *G. Wotawa replied that the recent incident in North Korea was known of at his centre within two hours and predictive studies were undertaken. They receive only analyses from ECMWF: it takes 72 hours to process data from the time of collection after an event, so analysis data are sufficient for these purposes. If required, forecast data are obtained from the US.*

EUMETSAT

EUMETSAT

Dr Fritz G. Wollenweber – EUMETSAT

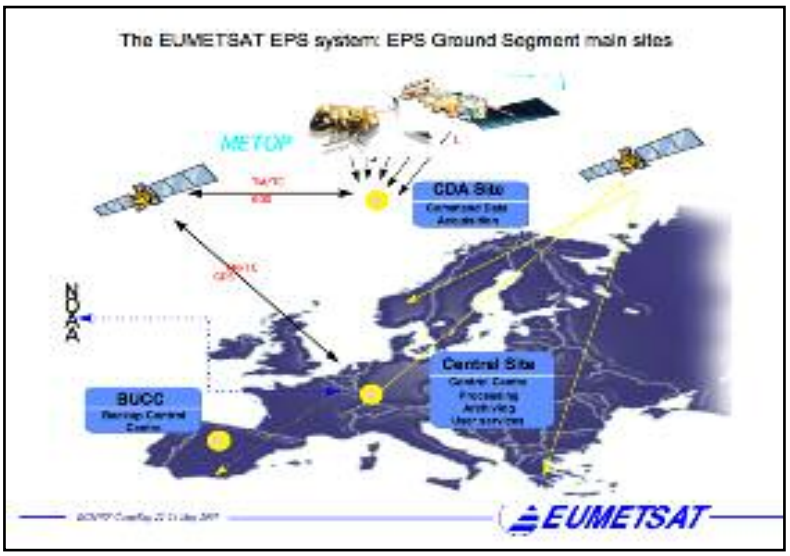

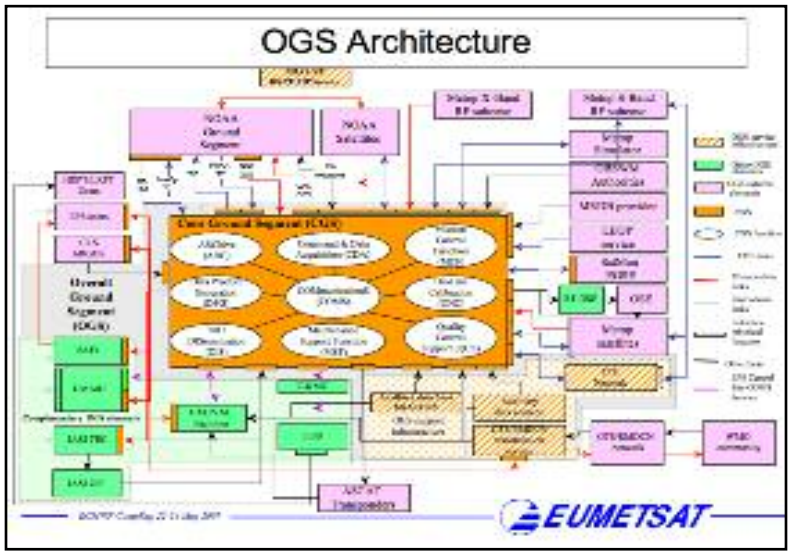

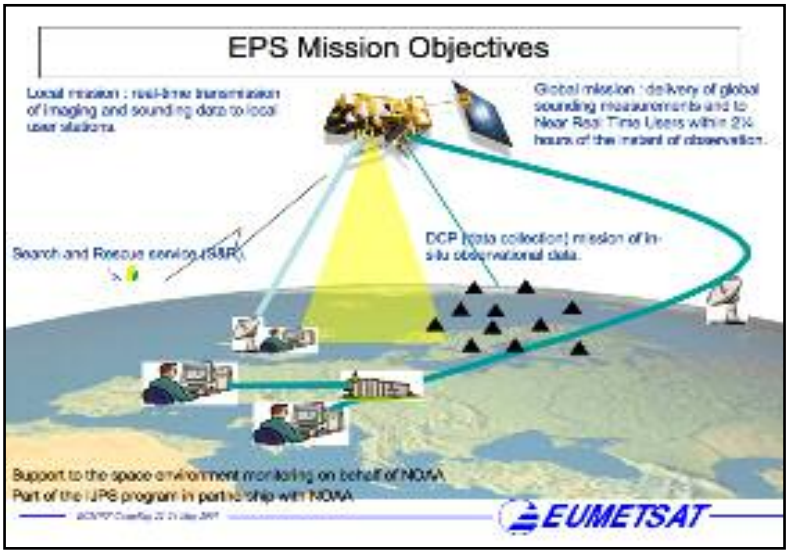

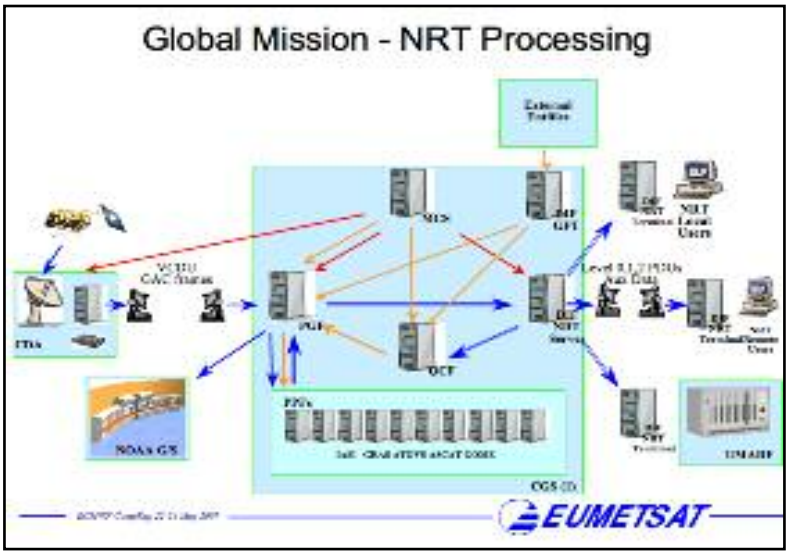
Global Mission - NRT Processing
EUMETSAT

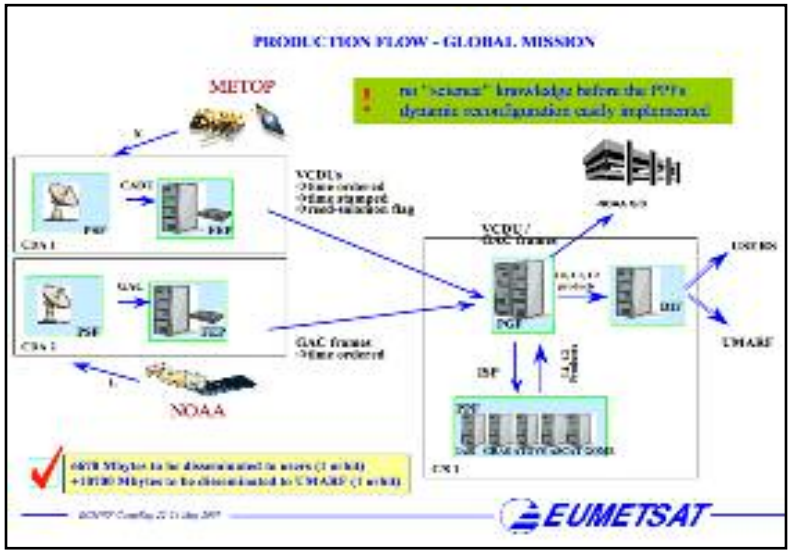
PRODUCTION FLOW - GLOBAL MISSION
METOP
NOAA
EUMETSAT

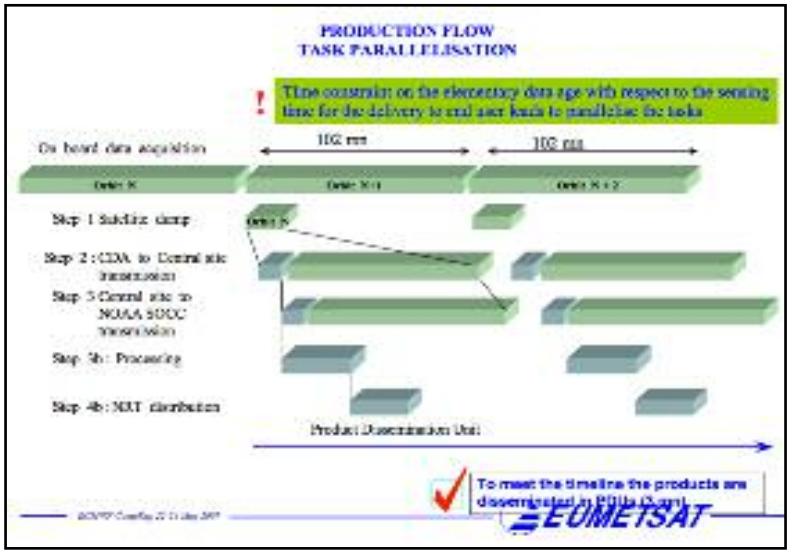
PRODUCTION FLOW
TASK PARALLELISATION
On board data acquisition
EUMETSAT

EUMETSAT **EUMETSAT**

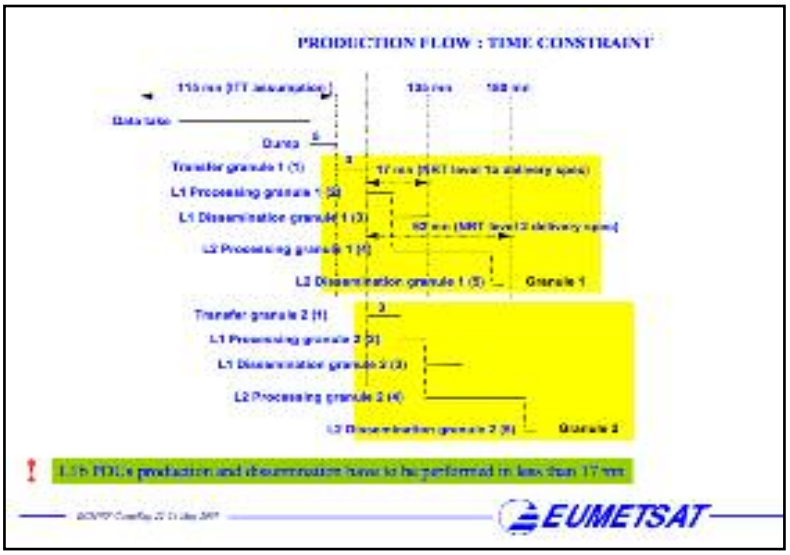

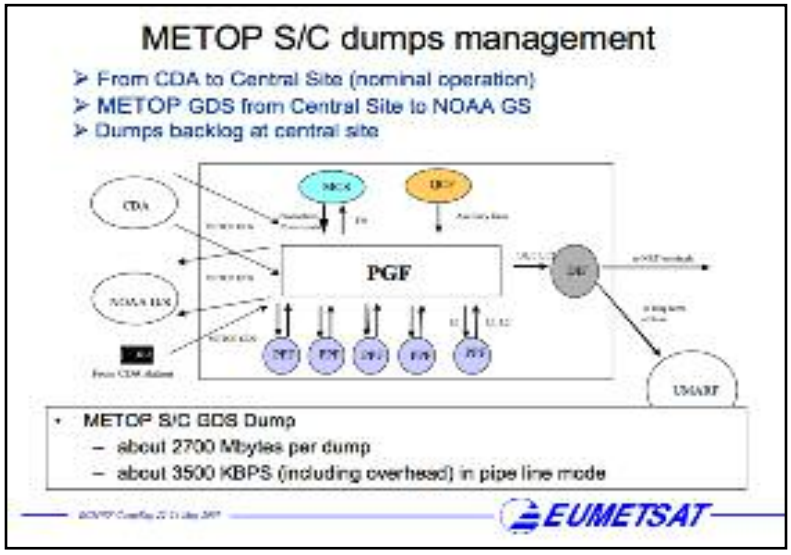

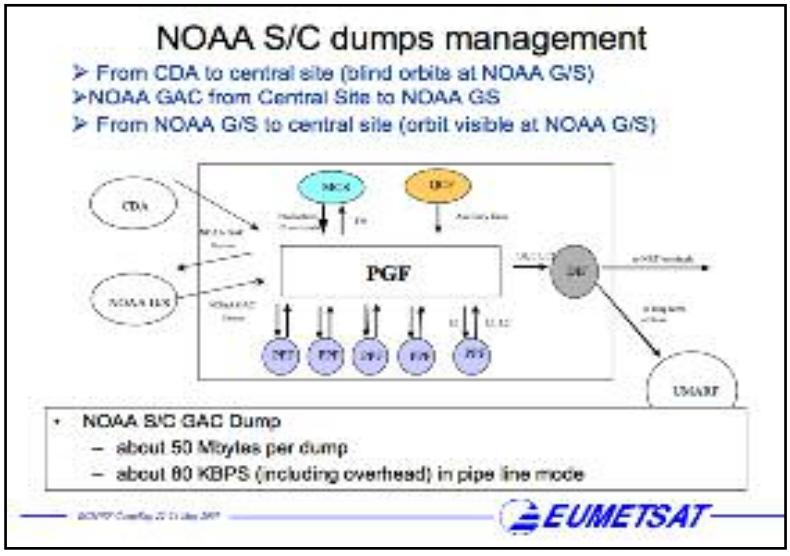

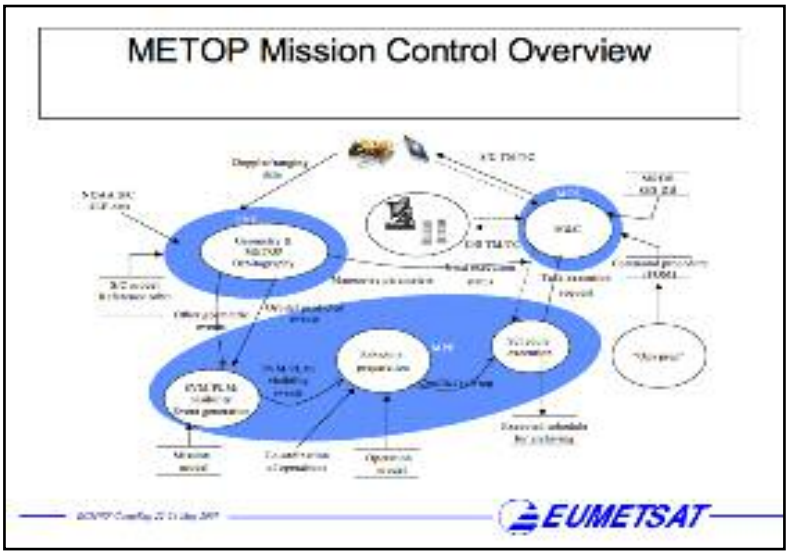
METOP Mission Control Overview
EUMETSAT

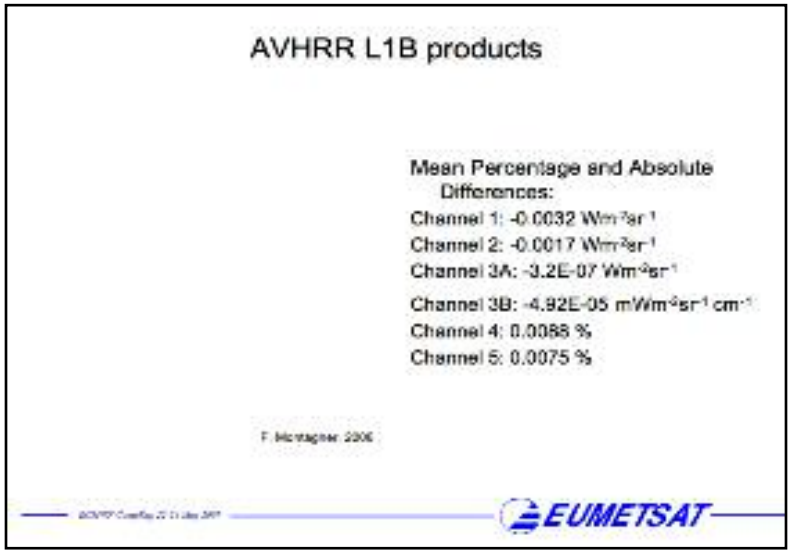
AVHRR L1B products
Mean Percentage and Absolute Differences:
Channel 1: -0.0032 Wm⁻²sr⁻¹
Channel 2: -0.0017 Wm⁻²sr⁻¹
Channel 3A: -3.2E-07 Wm⁻²sr⁻¹
Channel 3B: -4.92E-05 mWm⁻²sr⁻¹cm⁻¹
Channel 4: 0.0088 %
Channel 5: 0.0075 %
EUMETSAT

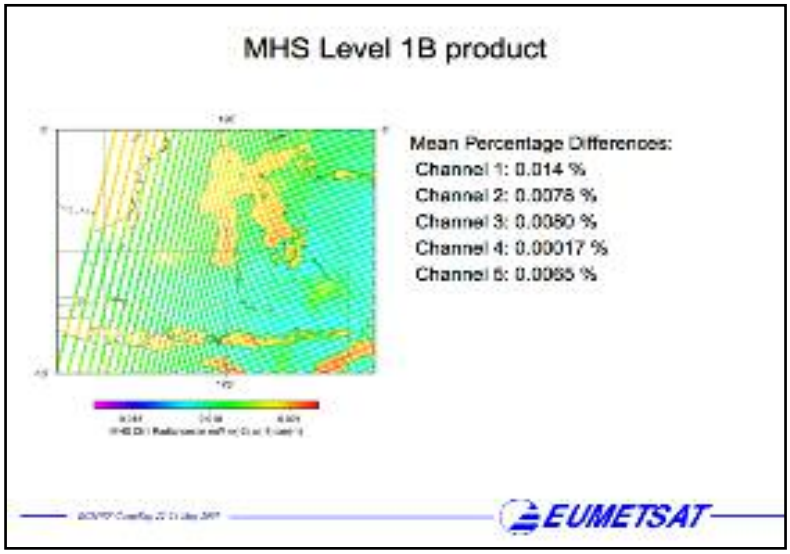
MHS Level 1B product
Mean Percentage Differences:
Channel 1: 0.014 %
Channel 2: 0.0078 %
Channel 3: 0.0080 %
Channel 4: 0.00017 %
Channel 5: 0.0065 %
EUMETSAT

EUMETSAT **EUMETSAT**

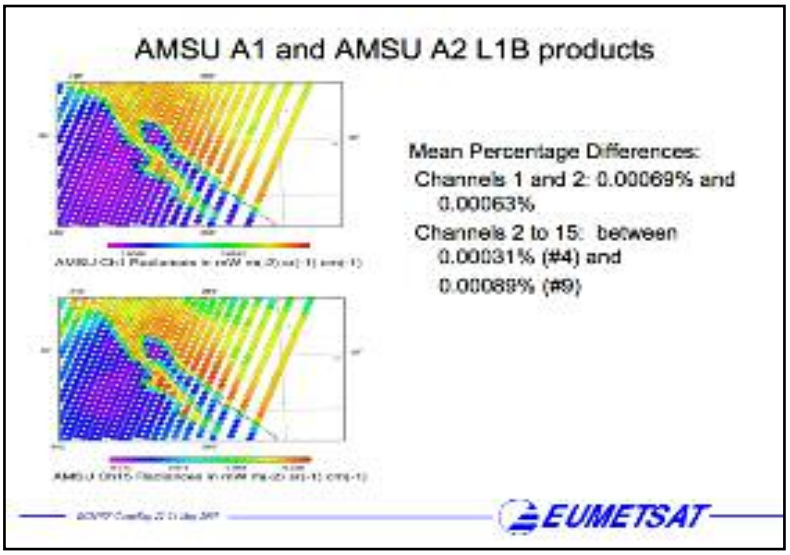

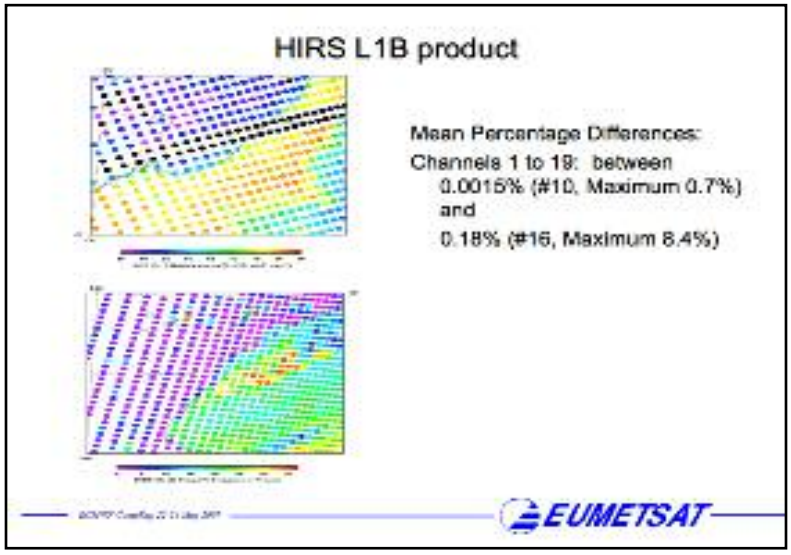

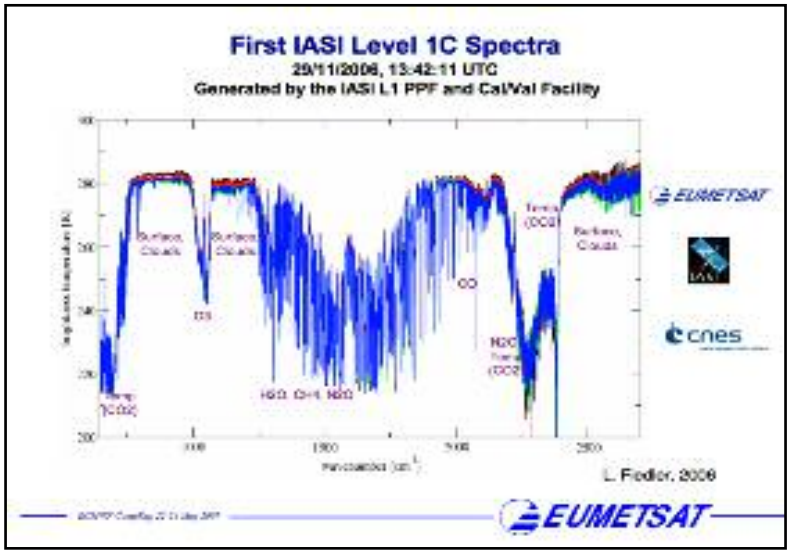

JRC

JRC

Stefano Galmarini – European Commission-DG-Joint Research Center / Institute for Environment and Sustainability

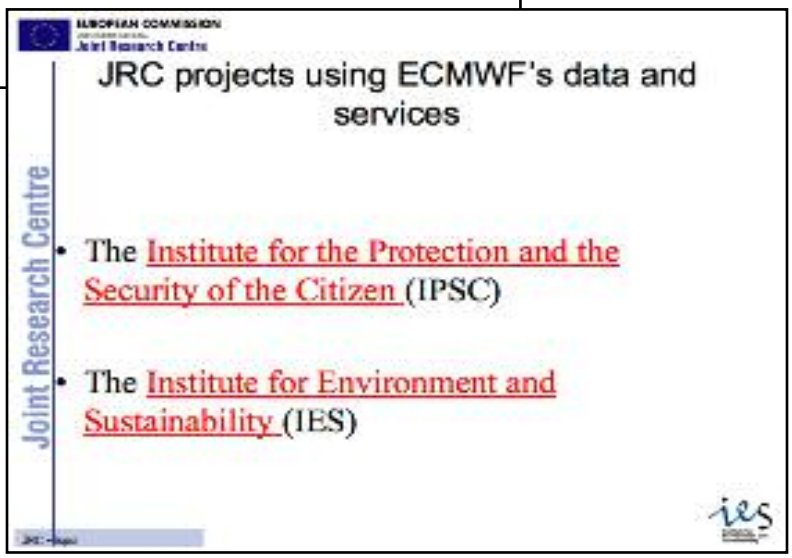

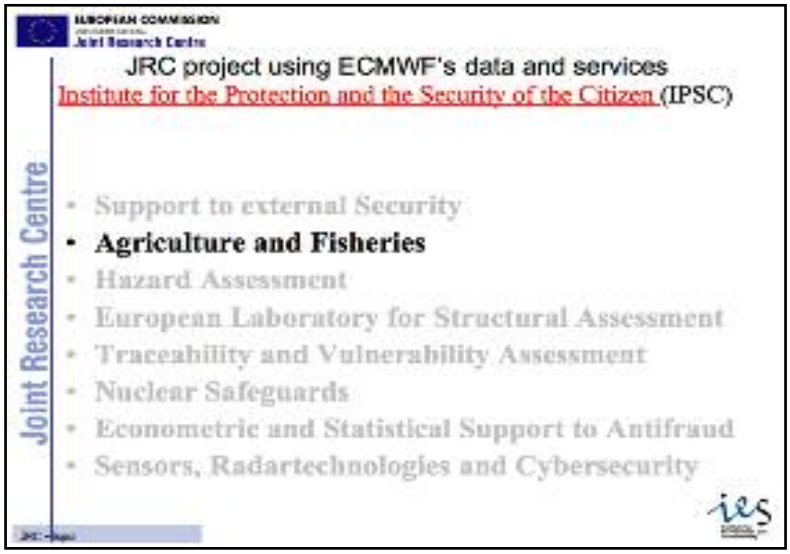

JRC JRC

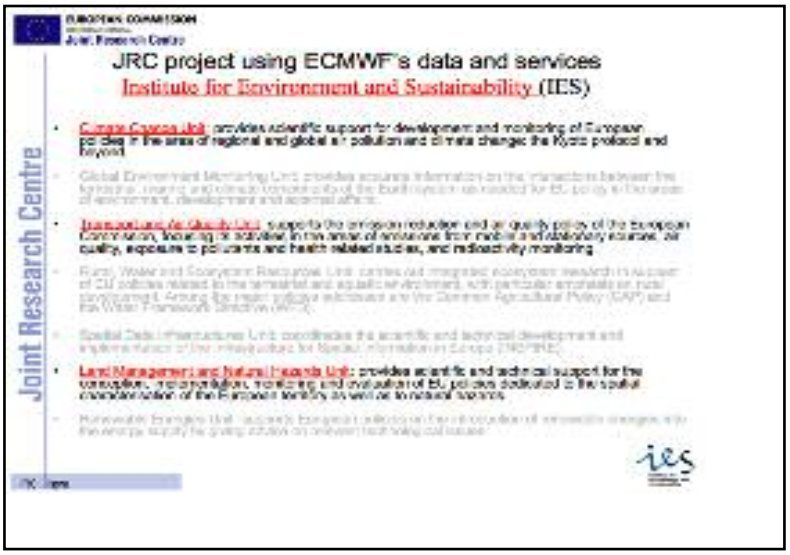
Joint Research Centre
JRC project using ECMWF's data and services
Institute for Environment and Sustainability (IES)
Land Management and Natural Hazards Unit: provides scientific and technical support for the


Joint Research Centre
Services in use
Agriculture and Fisheries
– ACCESS MARS ARCHIVE
Climate Change Unit:
– ACCESS MARS ARCHIVE
– USE OF HPCF
Transport and Air Quality Unit:
– ACCESS MARS ARCHIVE
Land Management and Natural Hazards Unit:
– ACCESS MARS ARCHIVE
– ACCESS TO ENSEMBLE FORECAST

ANNEX 1 ANNEX 1

Nineteenth Meeting of Computing Representatives
ECMWF, Shinfield Park, Reading, U.K., 22–24 May 2007
Participants

Austria	Mr. Gerhard Hermann
Belgium	Mr Rudi Swennen
Croatia	Mr. Vladimir Malović
Czech Republic	Mr. Karel Ostatnicky
Denmark	Mr. Thomas Lorenzen
Finland	Mr. Mikko Aalto
France	Mr. Dominique Birman
Germany	Dr. Elisabeth Krenzien
Greece	Major Antonis Emmanouil
Hungary	Mr. Istvan Ihász
Iceland	Mr Vigfús Gíslason
Ireland	Mr. Paul Halton
Morocco	Mrs. Fatima Driouech, Mr. Hassan Haddouch, Mr. Khalid Lahlal, Mr. Abderrazak Lemkhenter
Norway	Ms. Rebecca Rudsar
Romania	Mr. Roland Cotariu, Mr. Catalin Ostroveanu
Serbia	Mr. Vladimir Dimitrijevic
Slovenia	Mr. Petar Hitij
Spain	Mr. Jesus Gomez, Mr. Eduardo Monreal
Sweden	Mr. Mattias Andersson, Mr. Rafael Urrutia
Switzerland	Mr. Peter Roth
United Kingdom	Mr. Richard Barnes
EUMETSAT	Dr. Fritz Wollenweber
CTBTO	Mr. Gerhard Wotawa
JRC	Dr. Stefano Galmarini
ECMWF	Tony Bakker
	Paul Dando
	Francis Dequenne
	Richard Fisker
	Anne Fouilloux
	Enrico Fucile
	Laurent Gougeon
	Dominique Lucas
	Carsten Maass
	Umberto Modigliani
	Dieter Niebel
	Pam Prior
	Stephan Siemen
	Isabella Weger

ANNEX 2 ANNEX 2

Programme

Tuesday, 22 May 2007

14.00	Welcome	
	ECMWF's computer status and plans	I. Weger
15.15	Coffee	
15:45	New DRS and update on DHS	F. Dequenne
	Member State server: status and plans	R. Fisker
	New Strong Authentication system	D. Garcon
	LAN: overview and resiliency issues	D. Niebel
	RMDCN update	T. Bakker
17.30	Cocktails	

Wednesday, 23 May 2007

09.00	Member States' and Co-operating States' presentations	
10.30	Coffee	
11.00	General Discussion	
12.00	Visit of Computer Hall (optional)	
12.30	Lunch	
13.30	Member States' and Co-operating States' presentations	
15.30	Coffee	
16:00	Member States' and Co-operating States' presentations	
	Graphics update	S. Siemen
	GRIB API update and new GRIB tools	E. Fucile
	ECaccess: status and plans	L. Gougeon
	Member State time-critical applications	D. Lucas
19:30	Informal dinner	

Thursday, 24 May 2007

09.00	Member States' and Co-operating States' presentations	
10.30	Coffee	
11.00	New Incident Management System	P. Dando
	Usage of HPCF resources by Member States	U. Modigliani
	Final Discussion	
12.30	End of meeting	

ANNEX 3

ANNEX 3

Erwan Favennac – MétéoFrance