WORLD METEOROLOGICAL ORGANIZATION

COMMISSION FOR BASIC SYSTEMS

OPAG ON INFORMATION SYSTEMS & SERVICE

Expert Team on WIS-GTS Communication Techniques and Structure

Toulouse, France, 26-30 may 2008

ET-CTS/Doc. 6.3(2)
(16.V.2008)

ITEM 6.3

ENGLISH only

Lessons learned from the RA VI RMDCN & IMTN (Cloud II) migration to MPLS

(Submitted by ECMWF)

Summary and purpose of document

This document describes the migration process of the RMDCN from Frame-Relay to MPLS and the lessons learned.
This document is an updated version of the paper presented in November 2007 for the ICT-MTN meeting.

ACTION PROPOSED:

The meeting is invited to review the document.

TABLE OF CONTENT

1. Why migrating to MPLS? 2
2. From Frame-Relay to MPLS, a technical comparison 4
The Frame-Relay technology 4
The MPLS technology (Multi-Protocol Label Switching) 5
The new MPLS-based RMDCN network 5
3. Technical solutions 6
4. The Implementation 7
Validation of the new network 7
Migration of operational traffic 8
5. Performance and Reliability 8
6. Major incidents on the network during the last year 9
7. ECWMF Role and lessons learned 10

1. Why migrating to MPLS?

The RMDCN has been in operation since March 2000. It provides a network infrastructure for both the connections between ECMWF and its Member States and most of the GTS connections for WMO Regional Association VI. Over time it has expanded to encompass the Far East with connections to Japan, China and India. Currently there are 42 User Sites connected to it. Orange Business Services (OBS), formerly known as EQUANT, is the provider of this network. ECMWF manages the project and monitors the network on behalf of the connected User Sites following an agreement with WMO.

The initial network was based on a Frame Relay infrastructure using Permanent Virtual Circuits (PVC) between User Sites. There have been regular reviews of the contract which concentrated on pricing issues and also looked at the technology used for the network infrastructure. It was found that networks using Multi Protocol Label Switching (MPLS) were becoming more and more the norm. The RMDCN Operations Committee (ROC) and the ECMWF Technical Advisory Committee (TAC) were regularly updated on the progress of these types of networks, both in terms of reach and reliability and in terms of commercial viability. In 2004 OBS made an offer for a migration of the RMDCN from Frame Relay to an MPLS-based IPVPN (Internet Protocol Virtual Private Network) solution. The main features of the offer from OBS were as follows.

- A doubling of the bandwidth for the network connection for the same charge.
- An improved backup service with the effect that OBS were able to offer improved availability figures (the majority of sites now have 99.9% availability, while Mission Critical sites have 100%). A pre-condition for these improved figures is that the access circuits for both the primary and the backup connection are diversely routed from User Site to OBS Point of Presence.
- The provision of Class of Service (CoS) to allow traffic prioritization.
- Any-to-any connectivity.
- The ability of sites to choose a reduced service type (Silver service versus Gold).

However some technical trade-offs were made.

- Frame Relay networks provide site-to-site bandwidth guarantees by default, while MPLS-based networks only provide a bandwidth guarantee on the access into the MPLS distributed router (referred to as the cloud).
- In a Frame Relay network the traffic is automatically partitioned by the PVC infrastructure, whereas in an MPLS network the infrastructure is a shared resource and therefore allows for better utilization of the available bandwidth.
- The CoS feature in MPLS networks allows for traffic prioritization and substitutes the end-to-end bandwidth guarantees. Also the network management overhead in a Frame Relay network is quite significant. For an MPLS network this is greatly simplified.

Figure 1 and 2 shows the differences between the two network infrastructures, before and after the migration to the MPLS infrastructure.

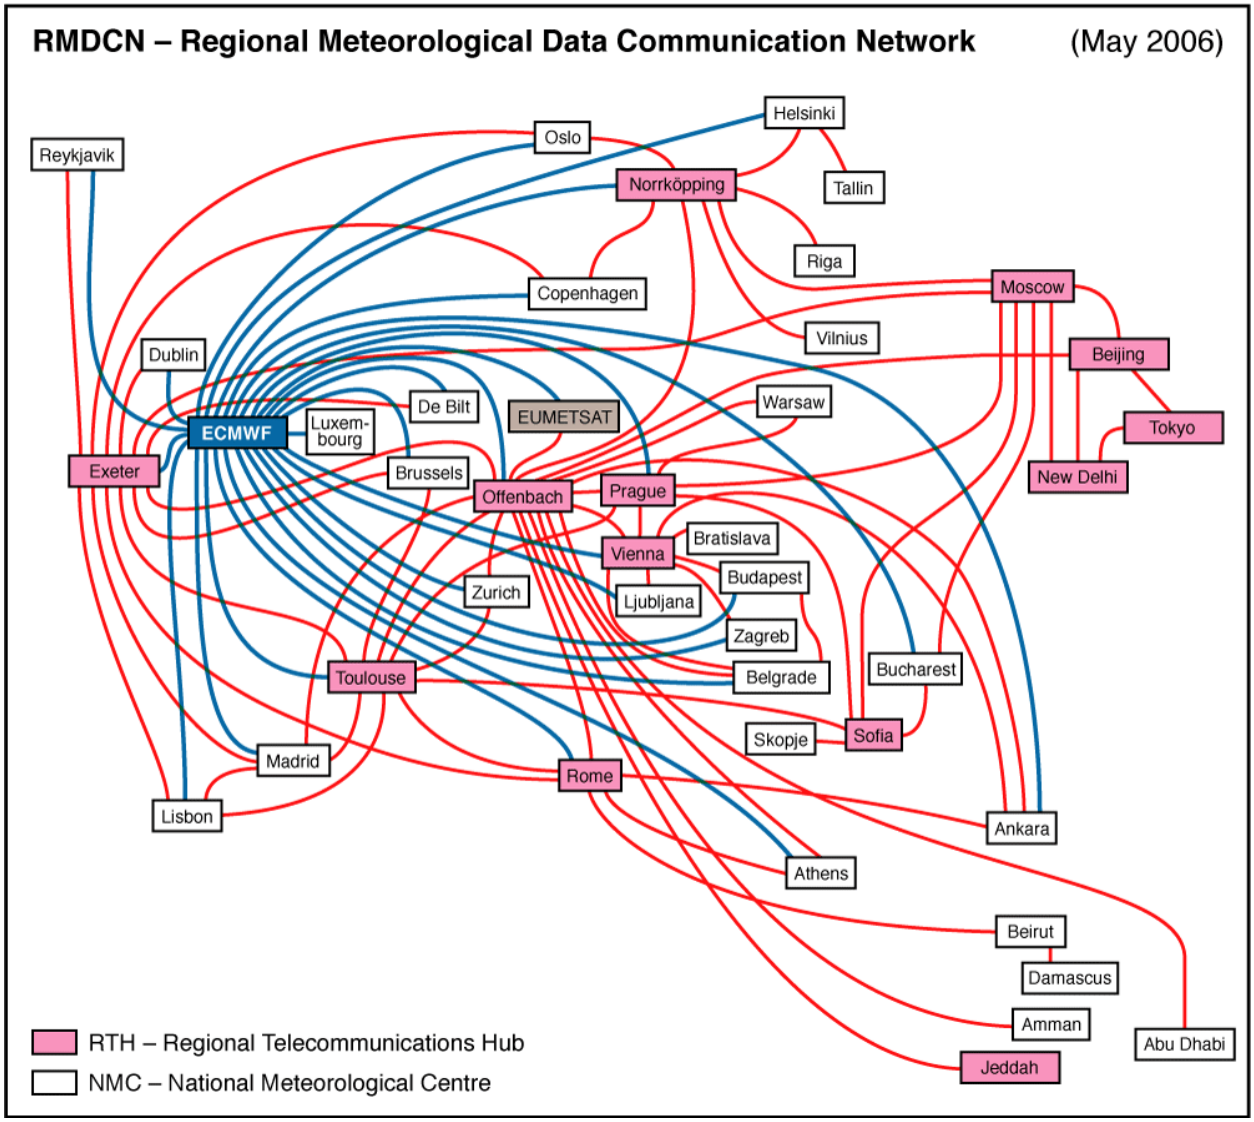

Figure 1 - Frame Relay based RMDCN

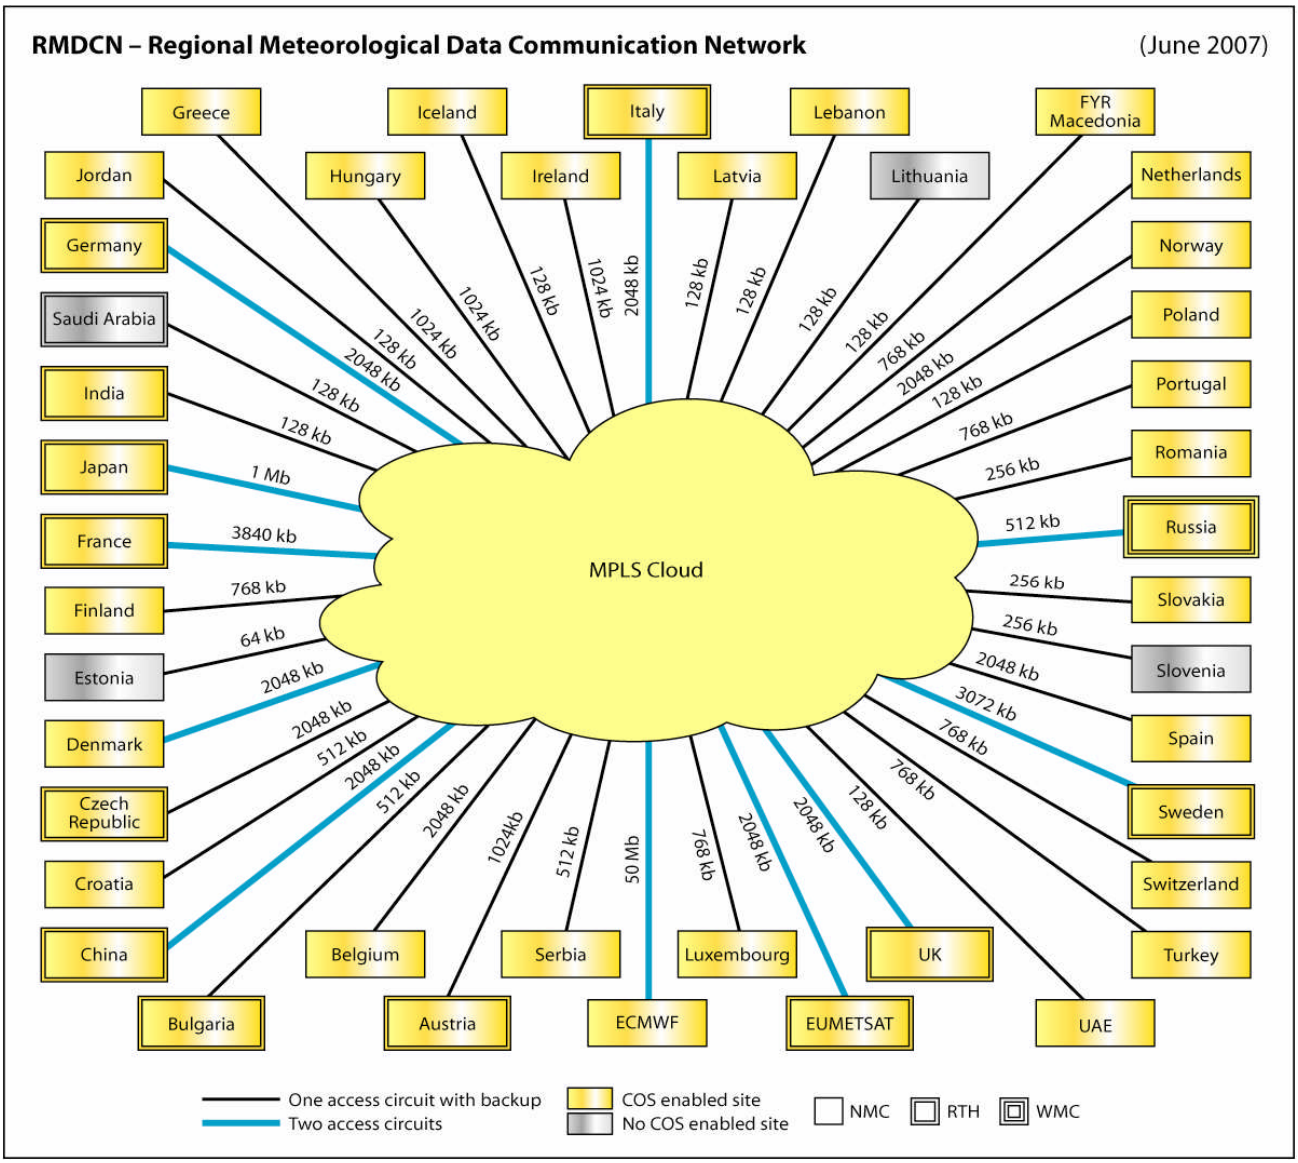

Figure 2 - MPLS Based RMDCN (May 2007)

The proposal from OBS for the migration of the network was approved by the ECMWF Council in December 2004 and this was also accepted by the other WMO members connected to the RMDCN. Since the new network infrastructure was significantly different the Service Level Agreement had to be revised. Also the implementation plan required detailed discussion with OBS. In order to guarantee an uninterrupted service on the existing Frame Relay network it was decided to implement the new network in parallel with the old network. Following intense negotiations Supplement 4 to the RMDCN contract was signed on 8 May 2006.

The new MPLS based network went operational on 18 June 2007.

2. From Frame-Relay to MPLS, a technical comparison

The MPLS technology used by the new RMDCN network is different from that of the previous Frame Relay technology. We now consider the differences and the improvements that are provided by the new network.

The Frame-Relay technology

Frame Relay is a relatively old technology which has the following two major characteristics.

- *Link-based Network*: In order to be able to exchange data, two sites must share a common dedicated point-to-point circuit. This circuit is deployed by the provider and is called a Permanent Virtual Circuit (PVC). Typically, a site shares several PVCs with all the sites with which it needs to exchange traffic.

- *Dedicated bandwidth between two sites*: When a PVC is deployed between two sites, it provides a dedicated bandwidth. This also highlights a limitation, as unused PVC bandwidth can not be (re)used by other PVCs.

For example Figure 3 shows a simplified Frame Relay network. Note that sites B and D cannot exchange data as they do not share a PVC.

The MPLS technology (Multi-Protocol Label Switching)

MPLS is an IP-centric solution that includes significant new features such as CoS (Class of Service) and VPN (Virtual Private Network). It is “Multi Protocol” as it can utilise different network technologies (FR, Ethernet etc.). It performs "Label Switching" since the data packets are switched through the network by virtue of an attached label. The following features characterize an MPLS-based solution.

- *One access circuit per site*: In order to exchange data with any of the other sites that are connected to the network, a site needs only one access circuit to the MPLS "cloud".

- *Dedicated bandwidth to the network*: The bandwidth value of the access circuit is the only guaranteed bandwidth for a site.

- *Traffic prioritisation and CoS (Class of Service)*: Through its traffic engineering mechanisms, MPLS provides a granular way of distinguishing the different traffic flows that cross the network and assign them the appropriate priority. Critical traffic can therefore be allocated a higher bandwidth.

Typically with MPLS a site can exchange data with all of the sites that are connected to the network. In essence, the network acts as a private Internet-like topology. This is known as any-to-any setup (see Figure 4). The lack of dedicated bandwidth between sites is compensated by the CoS configuration that allocates a higher priority the more critical the traffic is.

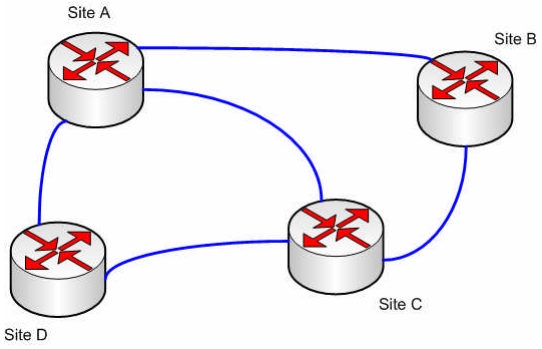

Figure 3

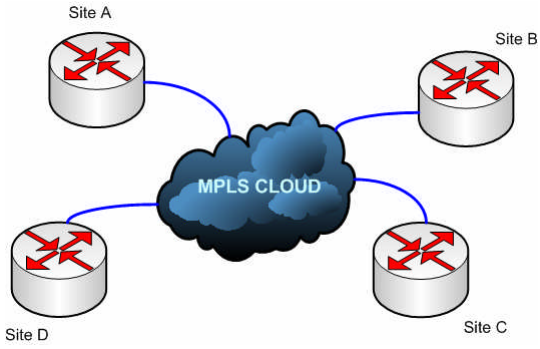

Figure 4

The new MPLS-based RMDCN network

For the RMDCN community, the two main benefits of an MPLS-based RMDCN network over one based on Frame Relay are the following.

- *Any-to-any connectivity*: Without further changes to the existing network, any pair of RMDCN sites can exchange data as soon as they have implemented the necessary routing changes locally.

- *Class of Service*: Taking into account that there is no dedicated bandwidth between the different RMDCN sites, the emphasis is on the traffic prioritisation and CoS. It is important that each RMDCN site classifies its traffic properly.

3. Technical solutions

During the specification phase of the project each site had to decide what kind of connection was needed.

For the connection, OBS offered:
- Mission Critical : Dual connection (to two different POP), Dual CPE (the second is a backup)
- Extra enhanced : Dual Connection (to two different POP), One CPE
- Enhanced : Dual CPE, one with ISDN access to another POP

For the Class of Service, OBS offered:

	D1	**D2**	**D3**
SILVER	N/A	100% of IP bandwidth	N/A
GOLD	75% of IP bandwidth	20% of IP bandwidth	5% of IP bandwidth

Note : The COS function only kicks in when there is congestion on the access line

4. The Implementation

The actual migration project commenced following the signature of Supplement 4 of the RMDCN contract on 8 May 2006.

Date	Task
8 May 2006	Signature of Supplement 4 of the RMDCN contract
11 September 2006	Start of installation – roll-out
16 April 2006	Ready for Trial Date – OBS handed over the network
	Start of the User Site Acceptance Tests
19 May 2007	Start of the Reliability Acceptance Test
4 June 2007	Start of the migration of all operational traffic
18 June 2007	Final acceptance of the new network

Table 1 - The major milestones during the migration project.

The following phases for the implementation were identified.

Specification of the configuration for each User Site

During this period the final configuration details such as access speed, mission critical setup, backup method, etc. were agreed with all RMDCN Members and the information was passed on to OBS.

Installation of the network

OBS rolled out the new MPLS-based network in parallel with the existing Frame Relay network. ECMWF gathered from all of the RMDCN members the detailed technical information (CoS classification, IP network addresses etc.) which OBS required to configure the network equipment.

Validation of the new network

In order to be confident that the new network was able to meet the operational requirements it had to be validated. This validation was split into two parts.

Firstly there would be a two week User Site Acceptance period during which tests were run to identify any configuration or deployment issues for each User Site on the network. This phase was also used to validate whether or not the traffic on the network was properly classified according to the CoS specifications and that the expected performance levels could be attained. To achieve this, all User Sites were provided with a common set of test software that had been developed by ECMWF.

The second phase of the validation was the Reliability Acceptance Test. Before running any real operational traffic on the network it was necessary to validate the service provided by OBS. During this phase ECMWF duplicated its dissemination traffic that was using the operational network and transmitted this duplicated data over the new network to systems that had been set up for testing by some member sites. Some Regional Telecommunications Hubs did similar tests by transmitting a copy of their operational GTS traffic over the new network to some of their partner sites.

Migration of operational traffic

In order to migrate the 91 PVCs in a controlled manner a detailed time schedule for migrating the operational traffic flows (i.e. PVC) was agreed by all RMDCN Members.

5. Performance and Reliability

The figure 5 shows the performance improvement for ECMWF dissemination to a particular country between the Frame Relay and the MPLS network.

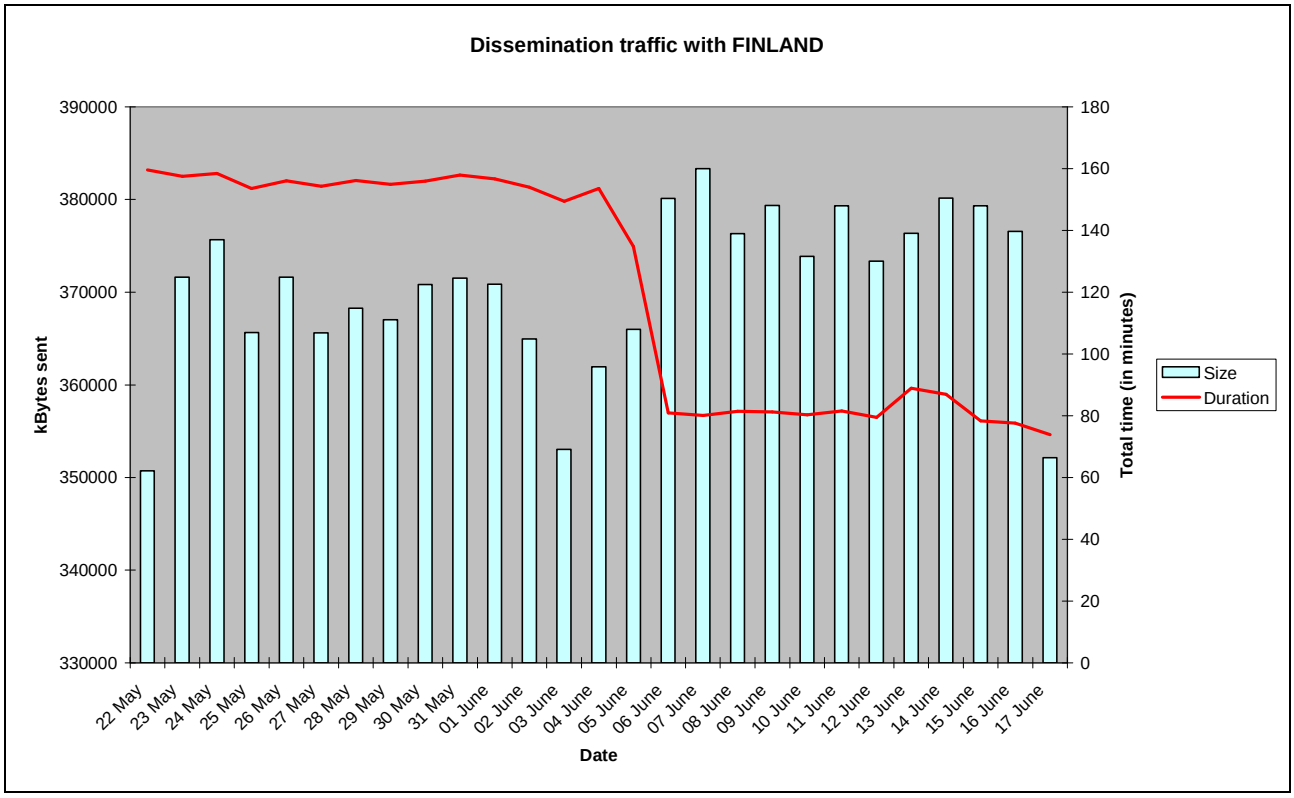

Figure 5: *Example showing the reduction in transfer times for ECMWF dissemination products. The blue bars denote the amount of data transferred (in kBytes). The red line denotes the duration of the dissemination (in minutes).*

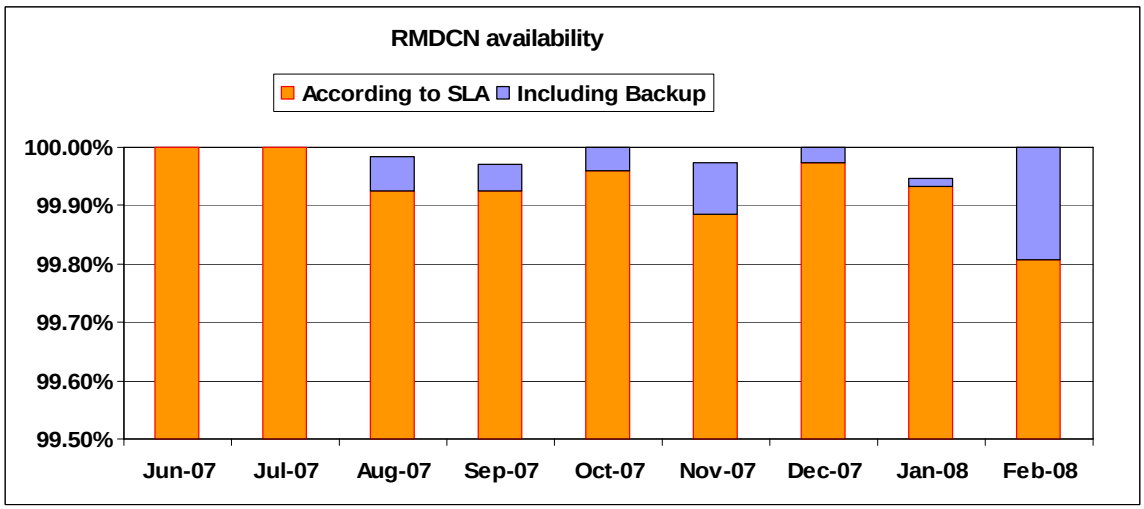

Figure 6: *RMDCN availability (according to SLA, and including backup)*

The pictures above prove that the network delivers the expected performance and meet the Service Level Agreement over the last months.

6. Major incidents on the network during the last year

- Italy - 8-9 Aug 2007: service down for 17.5 hours

  Problem in PTT exchange (access lines not completely diversely routed).

  Help Desk Feedback: “Access lines not on 24*7 support by local PTT”

- ECMWF-UK connection problem August 2007 – no impact on service

  BGP problem in OBS backbone

- Norway - 11-12 Sep 2007: service on backup for 15 hours

  Help Desk Feedback: “Access line not on 24*7 support by local PTT”

- Belgium - 2 Oct 2007: service down for 6 hours

  Cable problem at OBS POP

  ISDN not connected so no backup

- France – Bulgaria - November 2007: 2 hours down time

  OBS backbone issue (in reality problem was there for longer period but it was not reported in time and also workaround was found using France backup router)

- November 2007: Saudi Arabia, India and Romania suffered local PTT problems (SA does not have backup)

- Serbia – 15 jan 2008: 8 hours downtime; 30 Jan: 90 mins downtime

  Local PTT problem; no backup

- Portugal – 31 Jan 2008: 7 hours downtime primary connection – Backup ok

- OBS network problem - 13 Feb: affected sites UK, Iceland and ECMWF

  Equipment failure in OBS Backbone – backup worked for all sites

- Romania – 14 Feb: 5 hours outage primary connection

  Unscheduled maintenance – backup worked ok

These incidents triggered two actions for OBS

- Audit of diversity of Access Line for Mission Critical sites
- Verify 24*7 support local PTT

The overall performance and reliability are good.

7. ECWMF Role and lessons learned

As already mentioned ECMWF manages the project and monitors the network on behalf of the connected User Sites following an agreement with WMO.

This type of organization is rather uncommon for the provider. The final decision on what to do and how to do it is taken by the connected User Sites and not by ECMWF. OBS, and it is probably the same for others Telcos, is used to deal with one customer Centre responsible for all the connected sites. This is not our case, and this has proved to be a rather difficult situation to manage for OBS.

For such a large network, with so many different countries (very different level of technical knowledge, language barriers, number of staff involved…), ECMWF's role can be seen as:

- presenting a "unified" interface to the provider. For this to be successful, and the way the migration was done is a proof of that success, two key factors can be identified: a very good technical knowledge of networking, and a deep understanding of the meteorological community. Thanks to its position and to the people involved, ECMWF can provide both.

- supporting the connected countries with the daily activity of the network and with changes. This, for example, covers: providing statistics on the links, supervising the network 24x7, following the trouble tickets, checking SLA, checking invoices, dealing with modifications…

Obviously, such a migration is a major milestone in the life of the network. In particular, the change of the technology (from Frame Relay to MPLS) implied a lot of coordination between OBS, ECMWF and the User Sites. In order to minimize the risks and to keep the migration phase as short as possible, it was decided to use a "Big Bang" approach.

With this solution, the new network was built by the provider in parallel to the existing network (all routers configured, all leased lines in place…). Once everything was in place, the network was "delivered" to the community. During this phase, the Frame Relay network was still the operational network, OBS did not charge for the new infrastructure during that period. Though difficult to accept by OBS, it would have been impossible to proceed this way (a free of charge parallel network) if at the same time a different provider was used for the new network.

The "Big Bang" method proved its efficiency in terms of operation (almost no traffic disruption during the migration phase), however the 4.5 months delay in the project was related to this method. For the RMDCN migration this was not a major issue but might be one in some cases. We clearly set the priority on minimizing operational impacts and accepted the risk of delays.

Dealing with a lot of countries during this migration was really challenging in terms of the time schedule. On a day to day basis, there was obviously the issue of the time difference between countries. We also saw that there are always holidays somewhere! Summer break in Nordic countries starts as early as mid-June, May is a month of public holidays more or less everywhere, Easter is not at the same period, and even week-ends are not on the same days in some Islamic countries… So, scheduling operations and following multiple actions with all the countries was a real challenge.

However, the organisation established for this migration was quite scalable. Each analyst at ECMWF was the main contact point for a subset of countries and at the same time, each analyst was responsible for a particular subset of tasks (Project Management, Development of the assessment software, collection of configurations…). Should the RMDCN grow larger, this would easily be expandable to a larger team.

The RMDCN IPVPN MPLS Configuration.

Country/Site	Access Speed	IPVPN Port Speed	Resiliency	CoS	Load Balancing	NAS Backup Speed
ECMWF Member States						
Austria	2M	1M	enhanced	Gold	NO	512
Belgium	2M	2M	enhanced	Gold	NO	384
Denmark	2M	2M	**mission critical**	Gold	NO	N/A
Finland	2M	768	enhanced	Gold	NO	256
France	4M	3M	**mission critical**	Gold	NO	N/A
Germany	2M	2M	**mission critical**	Gold	NO	N/A
Greece	1M	768	enhanced	Gold	NO	384
Ireland	1M	1M	enhanced	Gold	NO	512
Italy	2M	2M	**mission critical**	Gold	NO	N/A
Luxembourg *	2M	768	enhanced	Gold	NO	N/A
Netherlands	2M	768	enhanced	Gold	NO	384
Norway	2M	2M	enhanced	Gold	NO	1M
Portugal	768	768	enhanced	Gold	NO	384
Spain	2M	2M	enhanced	Gold	NO	512
Sweden	4M	3M	**mission critical**	Gold	NO	N/A
Switzerland	2M	768	enhanced	Gold	NO	384
Turkey **	768	768	***extra enhanced***	Gold	NO	N/A
United Kingdom	2M	2M	**mission critical**	Gold	NO	N/A
ECMWF	50M	50M	**mission critical**	Gold	**YES**	N/A
ECMWF Co-operating States						
Croatia	512	512	enhanced	Gold	NO	256
Czech Republic	2M	2M	enhanced	Gold	NO	1M
Estonia	64	64	enhanced	**Silver**	NO	64
EUMETSAT	2M	2M	**mission critical**	Gold	NO	N/A
Hungary	1M	1M	enhanced	Gold	NO	256
Iceland	128	128	enhanced	Gold	NO	128
Lithuania	128	128	enhanced	**Silver**	NO	128
Romania	2M	256	enhanced	Gold	NO	128
Serbia	512	512	enhanced	Gold	NO	256
Slovenia	256	256	enhanced	Gold	NO	256
Other RMDCN Member States						
Bulgaria	512	512	enhanced	Gold	NO	128
China	2M	2M	**mission critical**	Gold	NO	N/A
India	128	128	enhanced	Gold	NO	128
Japan	1M	1M	**mission critical**	Gold	**YES**	N/A
Jordan	128	128	enhanced	Gold	NO	128
Latvia	128	128	enhanced	Gold	NO	128
Lebanon	128	128	enhanced	Gold	NO	128
FYR Macedonia	128	128	enhanced	Gold	NO	128
Poland	128	128	enhanced	Gold	NO	128
Russian Federation	512	512	**mission critical**	Gold	NO	N/A
Saudi Arabia *	512	128	enhanced	**Silver**	NO	N/A
Slovakia	256	256	enhanced	**Silver**	NO	128
United Arab Emirates	128	128	enhanced	Gold	NO	64

*) Luxembourg and Saudi Arabia do not have a secondary (backup) connection

**) Turkey has got a secondary connection with a speed of 384 kbps