IPSec-VPN as a backup for the RMDCN

ROC-14, 3-4 June 2008, Vienna

Submitted by ECMWF

Introduction

- Project presentation – Background
- Configuration
- Tests results
- Other technical solutions
- Conclusion – Recommendations – Next Steps

Project Presentation – Background

- **2002: IPSec feasibility study: ECMWF, Germany, Greece, France and the Netherlands**
  - Provides **guidelines and recommendations** for building secure connections over the Internet

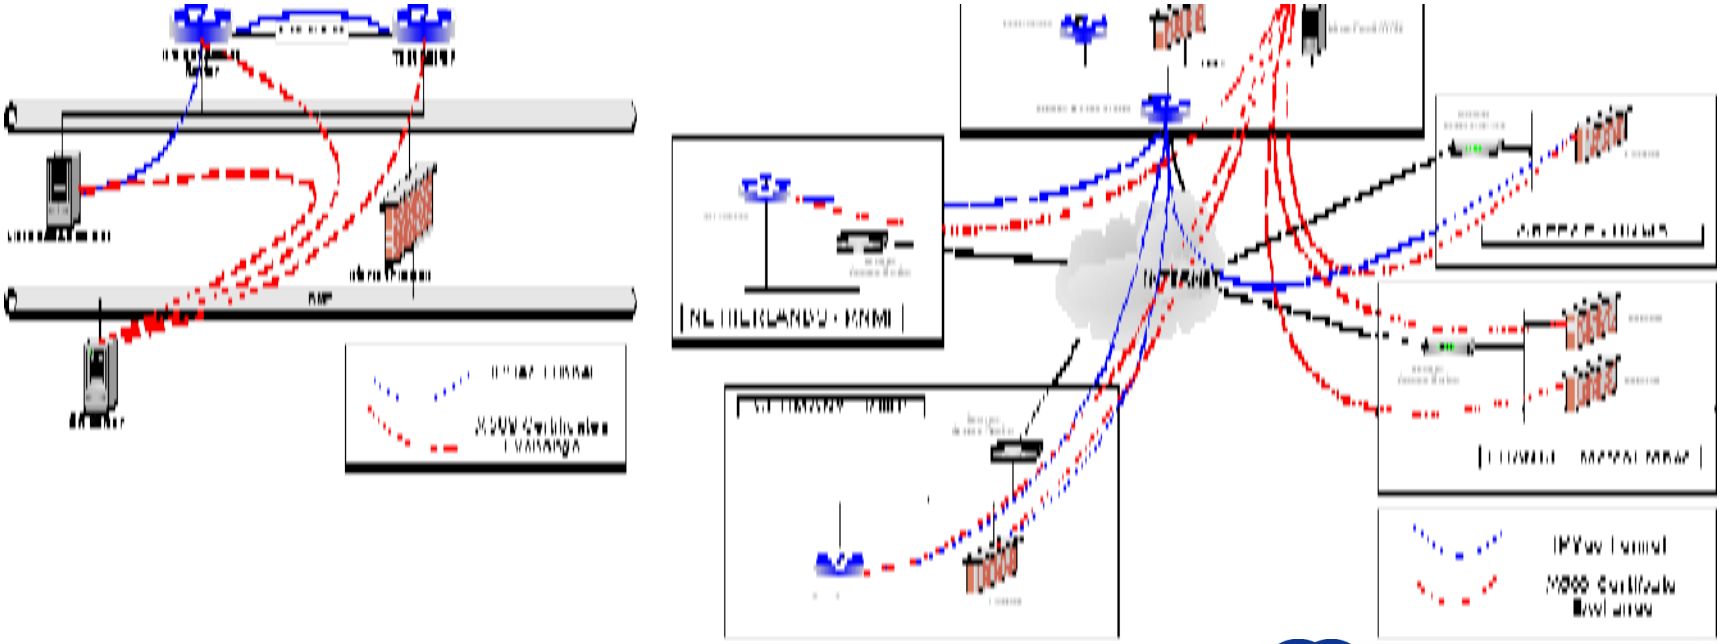

ECMWF

Project Presentation – Background

- **2005: IPSec-based VPN as a backup for the RMDCN study: ECMWF, Bulgaria, China, France, Germany, Sweden, the Russian Federation and the UK)**
  - Provides a **framework** for an operational RMDCN backup solution using an Internet-based IPSec VPN
  - Only "**static**" rerouting considered

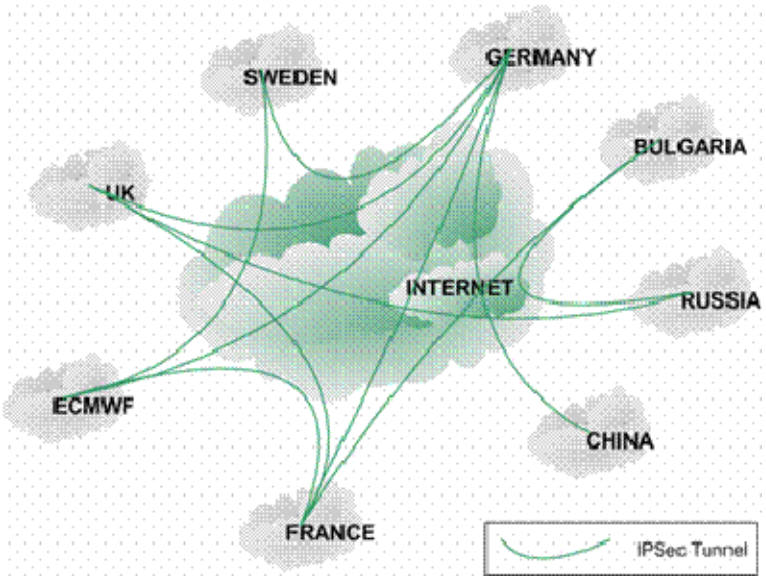

Project Presentation – Background

- **2007-2008: IPSec VPN Backup for the RMDCN project: ECMWF, Belgium, China, Germany and Turkey**
  - Using and IPSec-based VPN infrastructure to **transport operational RMDCN traffic between RMDCN sites** as an alternative to the RMDCN network itself
  - Phase #1: Building the IPSec-based infrastructure
  - Phase #2: Using the IPSec-based VPN infrastructure as a backup for the RMDCN in an operational context

Configuration

- **Inspired by OBS's RMDCN CAS failover implementation**
  - uses any-to-any connectivity, ECMWF is used as a "relay" between NAS domain and MPLS cloud.

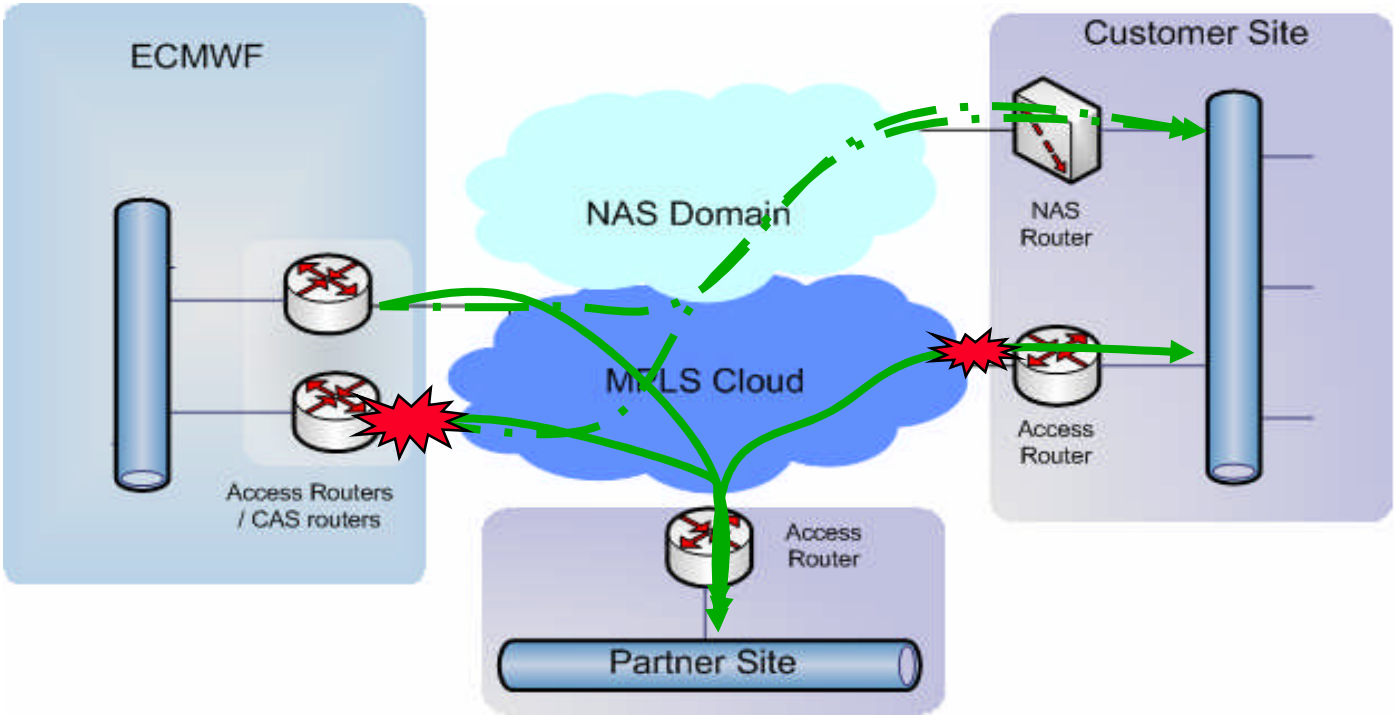

Configuration

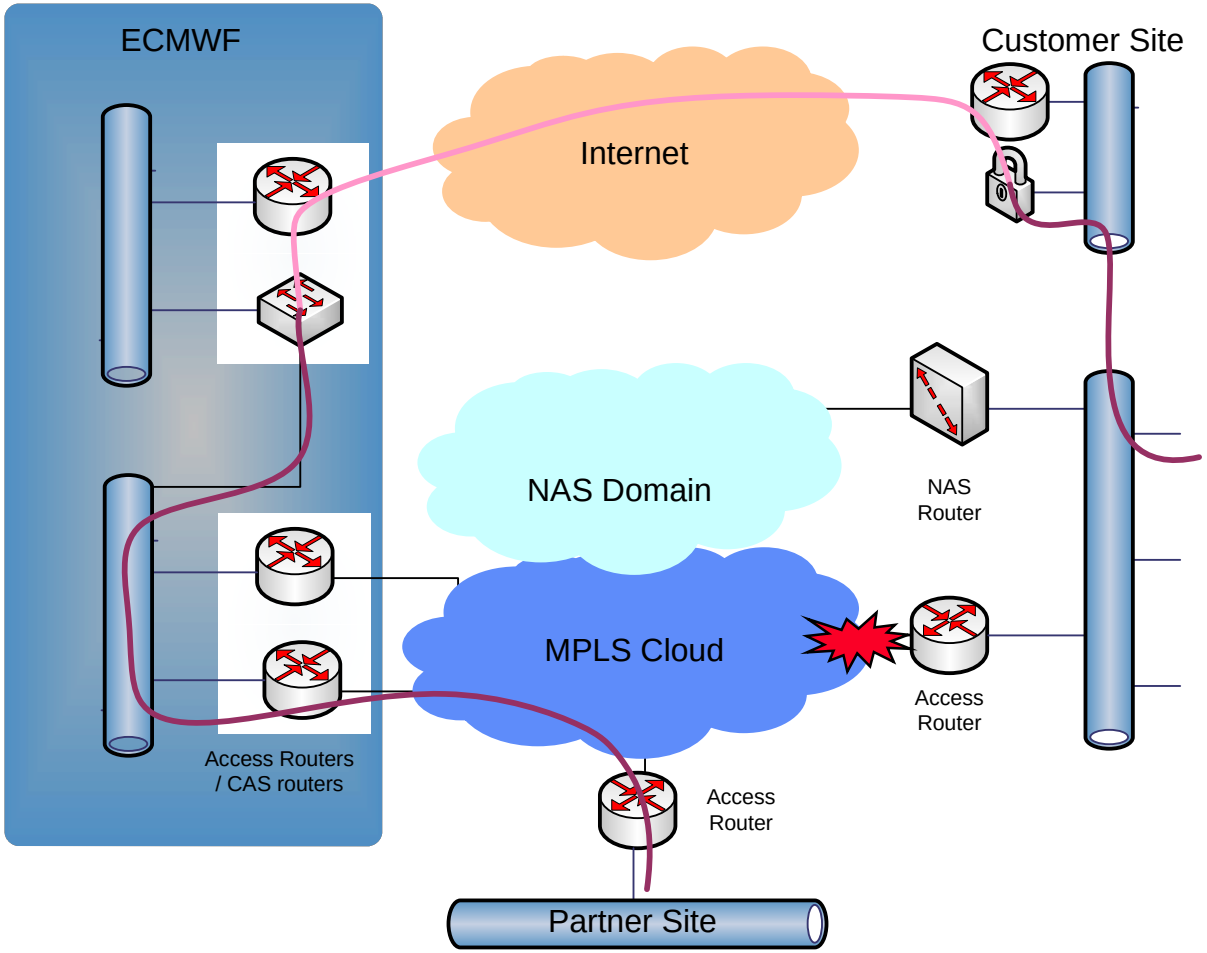

Configuration

- **Configuration at ECMWF**
  - Mimic the NAS ISDN backup implementation within the RMDCN: ECMWF acts as an IPSec centralising site, which guarantees the **any-to-any** connectivity of the RMDCN IPVPN cloud
  - ECMWF advertises the alternative routes to the RMDCN community through a dynamic routing exchange with the OBS IPVPN CE routers. The preferred routing protocol is **EIGRP**
  - OBS advertises the backup routes with a **lower priority** to the RMDCN IPVPN cloud through
    - Redistribution of the EIGRP routes into BGP
    - Implementation of a BGP-community tagging
  - OBS advertises the RMDCN IPVPN routes to ECMWF through the **redistribution** of the BGP routes into EIGRP

Configuration

- **Configuration at ECMWF**

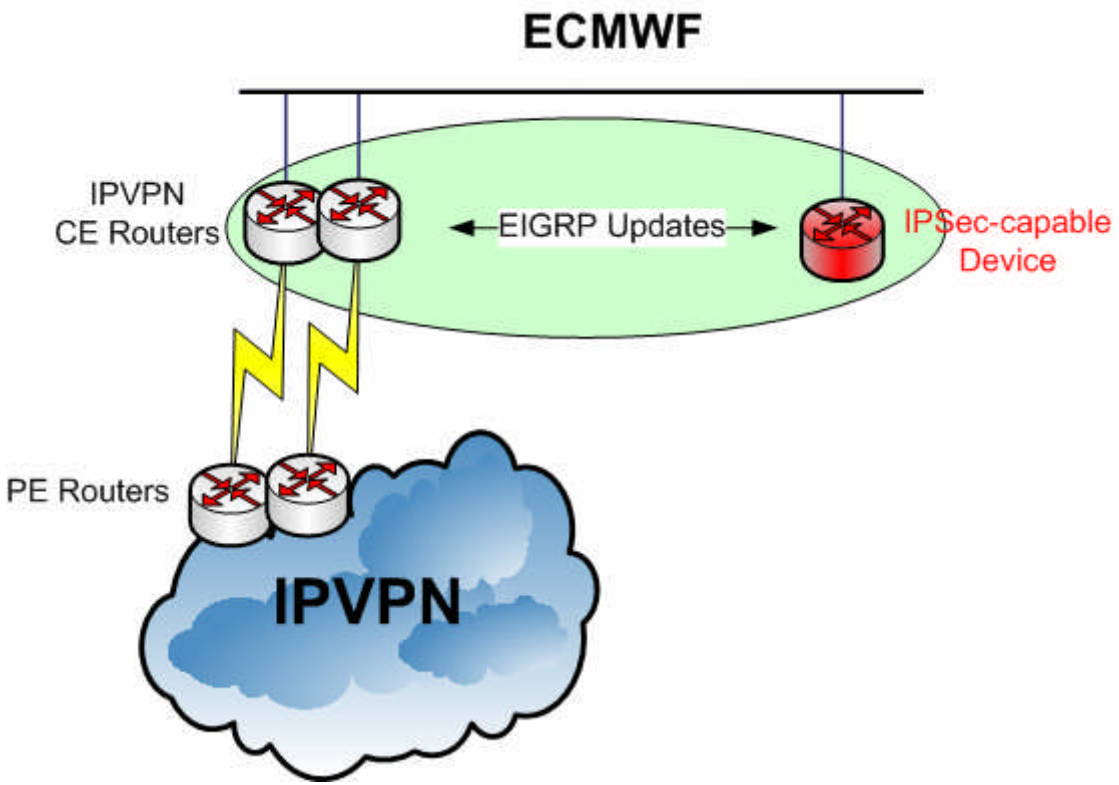

Configuration

- **Configuration at an RMDCN site – Option #1: using manual rerouting**
  - The IPVPN CE router is not involved in any dynamic protocol exchange with any local network equipment
  - In case of outage, the local site is responsible for manually rerouting the operational RMDCN traffic though the permanently established IPSec tunnel with ECMWF via the local site IPSec-capable device. This could be done either by using static routes or a dynamic routing protocol

Configuration

- ECMWF acts as proxy for the site:
  - It forwards the traffic to the RMDCN-IPVPN cloud
  - It receives any traffic sent towards this site through the activation of the backup route(s) using the EIGRP redistribution into BGP and forwards it through the IPSec tunnel

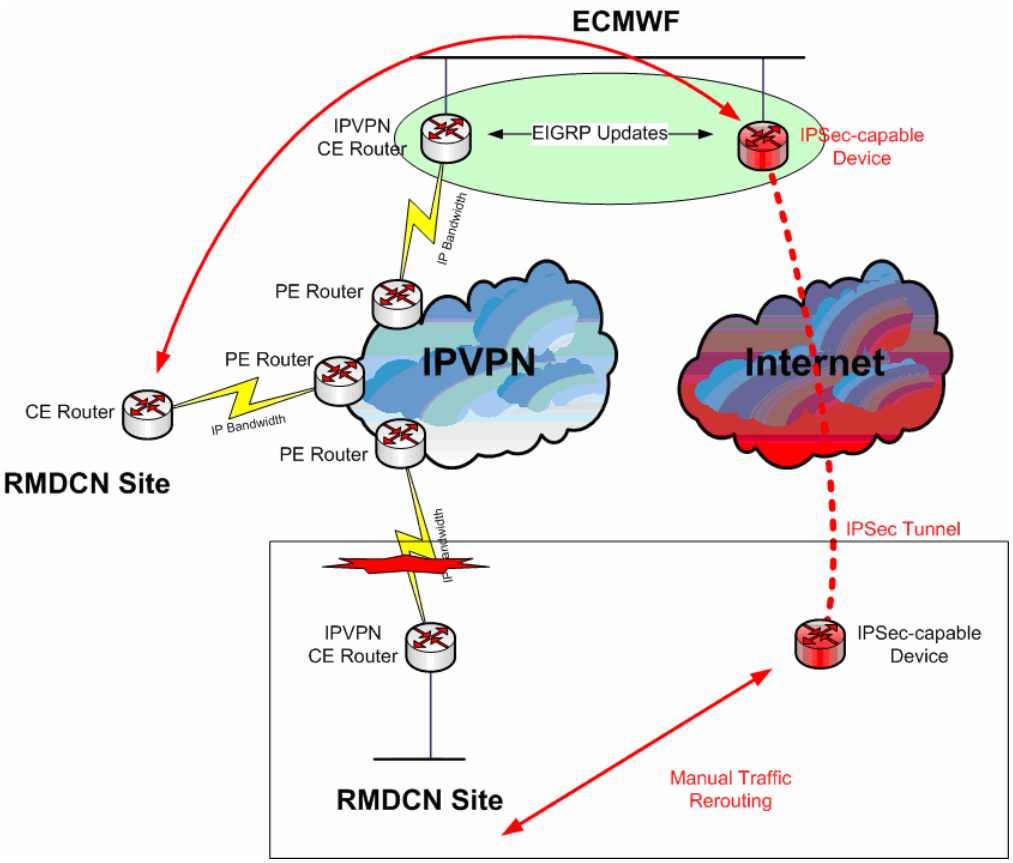

Configuration

- **Configuration at an RMDCN site – Option #2: using dynamic routing protocol**
  - A dynamic routing protocol exchange is implemented between the IPVPN CE router and a local network device
  - The local site uses this dynamic routing protocol to route its operational RMDCN traffic towards the IPVPN CE router
  - In case of outage, the local site network device reroutes all the operational RMDCN traffic towards the IPSec-capable device
  - The IPSec-capable device forwards this traffic to ECMWF through the permanently established IPSec tunnel with ECMWF

Configuration

- ECMWF acts as proxy for the site:
  - It forwards the traffic to the RMDCN-IPVPN cloud
  - It receives any traffic sent towards this site through the activation of the backup route(s) using the EIGRP redistribution into BGP and forwards it through the IPSec tunnel

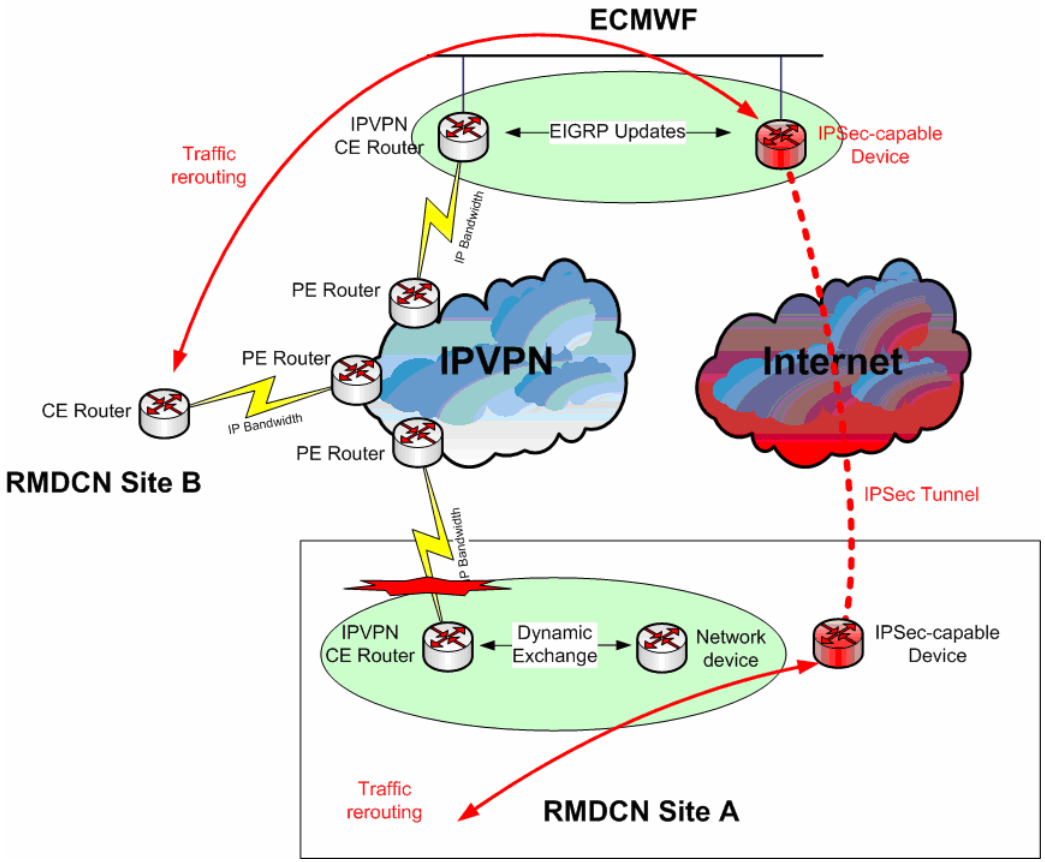

Configuration

- **IPSec tunnel specifications**
  - As recommended in the previous IPSec Studies:
    - Device authentication: X509 certificates, it is the most scalable and the most secure authentication method, using the already existing ECMWF Public Key Infrastructure (PKI)
    - IKE session key exchange: Deffie-Hellman group 2 or higher
    - Data integrity: the use of ESP HMAC, either MD5 or SHA
    - Data encryption: either ESP_NULL or ESP_AES
  - Tests preparations
    - Allow the IPSec traffic towards the remote sites
    - X509 certificate enrollment with the ECMWF CA server
    - Set-up of the IPSec configuration

Configuration

- **IPSec devices involved**

Country	IPSec device type	Hostname	OS version	IPSec device IP address	Test server IP address	RMDCN Networks
Belgium	Openswan (open source)	-	-	-	-	193.190.231.160/28 193.190.249.224/28
China	Cisco ASA	cmavpn	8.0(3)	210.73.54.50	57.206.141.144	57.206.141.128/26
ECMWF	Cisco ASA	ecvpn	8.0(3)	193.61.196.38	136.156.14.211	136.156.0.0/16
Germany	Nortel/Checkpoint	dwdfw	NGX R61	141.38.1.11	141.38.41.26	141.38.0.0/16
Turkey	Nokia/Checkpoint	MeteorCluster	NGX R65	212.175.180.4	57.206.143.220	57.206.143.192/26

Tests Results

- **Phase #1 – Building the IPSec-based VPN infrastructure**
  - The use of operational IPSec gateways: apart from China, all the sites were using operational IPSec device which meant that:
    - Each change had to be made very carefully
    - It was not possible to deploy a “final” IPSec configuration
  - All sites apart from China (EIGRP) use static routes to re-route the operational traffic through the IPSec tunnel (Option #1)
  - Checkpoint IPSec interoperability issues: establishing Cisco ASA to Checkpoint IPSec tunnels proved to be quite challenging (Turkey, Germany)
  - Nortel VPN accelerator card issue: in Germany, this card has to be disabled for the RMDCN traffic to go through the IPSec tunnel established with ECMWF

Tests Results

- **Phase #1 – Building the IPSec-based VPN infrastructure**
  - Results summary
    - Belgium: after considering using PIX, ipsec-tool and Openswan, the tests stalled and no working IPSec configuration could be implemented
    - China (ASA): building the IPSec tunnel was quite straightforward since both sites were using Cisco ASA devices
    - Germany (Checkpoint): establishing IPSec tunnels proved to be difficult, but everything was fine after disabling the Nortel VPN accelerator card and the configuration proved to be stable since
    - Turkey (Checkpoint): establishing IPSec tunnels proved to be difficult, although the configuration proved to be stable afterwards

Tests Results

- **Phase #2 – Using the IPSec-based VPN infrastructure as a backup for the RMDCN in an operational context**
  - Once OBS has activated the EIGRP redistribution for a site, a “live” re-routing was performed in 3 steps:
    1. Complete the IPSec tunnel configuration (if necessary)
    2. Simulate a link failure
    3. Revert back the changes (if necessary)

Tests Results

- **Phase #2 – Using the IPSec-based VPN infrastructure as a backup for the RMDCN in an operational context**
  - Results summary
    - China: test done 28/05. No steps #1 and #3 as **automatic** rerouting. The RMDCN traffic was **successfully** re-routed with no other delay than the BGP convergence time
    - Germany: test done on the 8th of May 2008. The **static** re-routing was **successful**. The RMDCN traffic towards 5 sites (out of 18) was not re-routed properly, but this is not strictly related to the IPSec infrastructure itself
    - Turkey: test done on the 16th of April 2008. The **static** re-routing with the three RMDCN sites that exchange data with Turkey was **successful** (ECMWF, Germany and Italy)

Other Technical Solutions

- **All Checkpoint – 2 Topologies**
  - "hub-and-spoke" topology ("**Star** VPN Community"): it is fairly easy to deploy a solution whereby all the IPSec traffic is centralised (at ECMWF for example) and each site encrypts its own traffic ("VPN domain"). This is similar to the deployed IPSec VPN configuration.

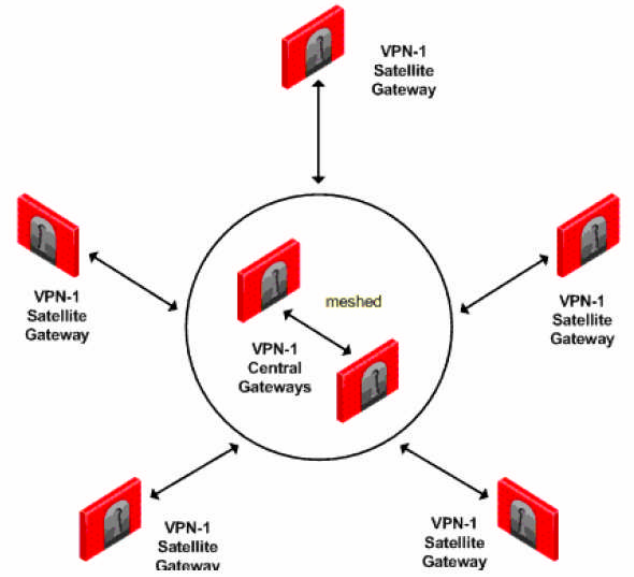

Other Technical Solutions

- **All Checkpoint – 2 Topologies**
  - "any-to-any" topology ("**Meshed** VPN Community"):
    - if all the gateways are **centrally managed**, this is easy to implement as the conf would be "pushed" to all the gateways
    - if, **more likely**, each gateway is locally managed, then we are back to a site-to-site conf and each site needs to manually define all the remote sites gateways

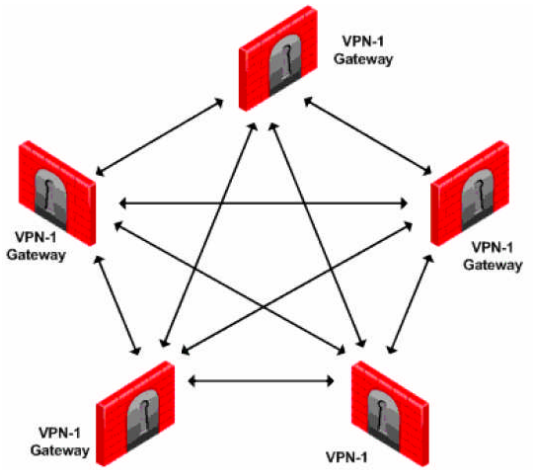

Other Technical Solutions

- **All Cisco – DMVPN**
  - This is a Cisco IOS solution for building IPsec+GRE VPNs
  - Relies on two proven Cisco technologies
    - Next Hop Resolution Protocol (NHRP): Hub maintains a (NHRP) database of all the spoke's real (public interface) addresses
      - Each spoke registers its real address when it boots
      - Spokes query NHRP database for real addresses of destination spokes to build direct tunnels
    - Multipoint GRE Tunnel Interface
      - Allows single GRE interface to support multiple IPsec tunnels
      - Simplifies size and complexity of configuration

Other Technical Solutions

- **Hub-and-spoke**
  - All VPN traffic must go via hub
    - Hub bandwidth and CPU utilization limit VPN
  - Number of tunnels = O(n)
- **Dynamic-Mesh – Dynamic spoke-spoke tunnels**
  - Control traffic — Hub to Hub and Hub and spoke
    - Next Hop Resolution Protocol (NHRP), Dynamic Routing, IP Multicast
  - Data traffic — Dynamic mesh
    - Spoke routers only need to support spoke-hub and spoke-spoke tunnels currently in use.
    - Hub only supports spoke-hub traffic and overflow from spoke-spoke traffic.
  - Number of tunnels > O(n), << O(n2)

Other Technical Solutions

- NHRP Resolution – Process Switching

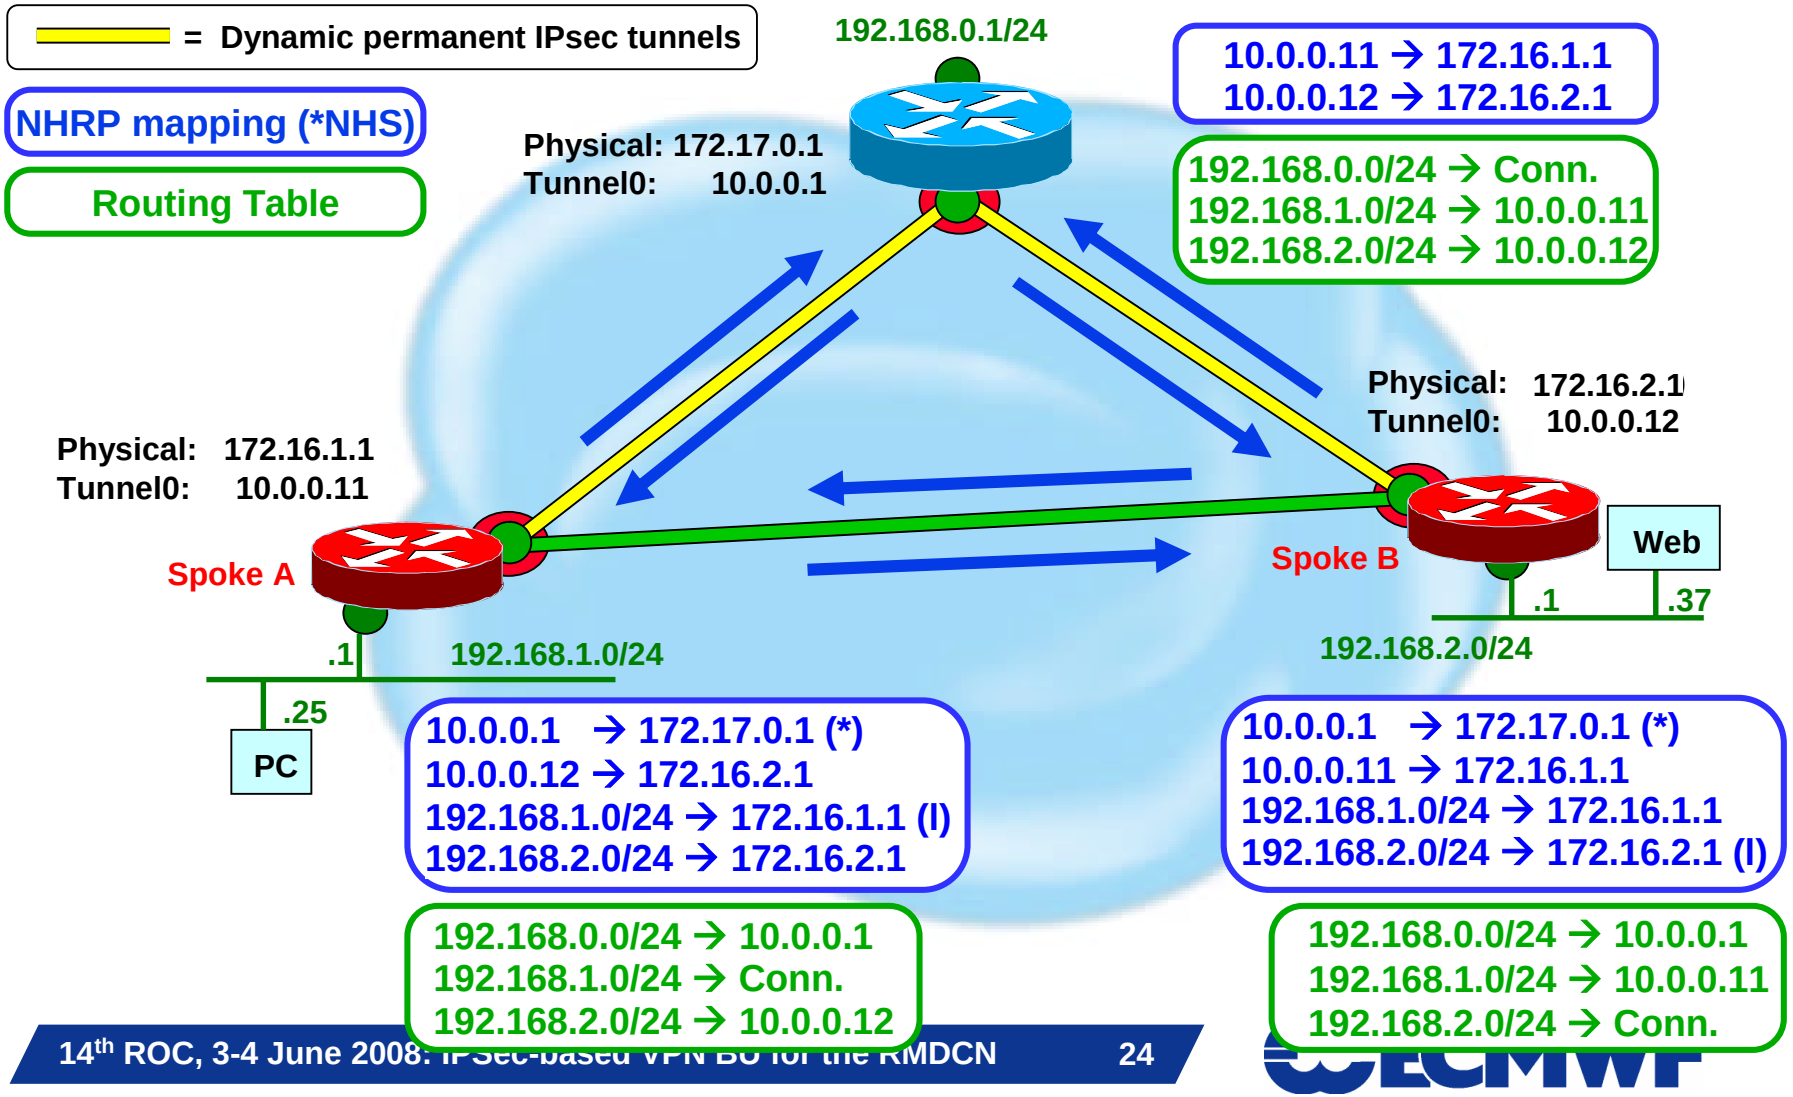

Other Technical Solutions

- **Summary**
  - All Checkpoint
    - Solution is more suitable for a centralised "Corporate" deployment
  - All Cisco – DMVPN
    - It does not alter the standards-based IPsec VPN tunnels, but it changes their configuration
    - Very scalable and easy to configure

Conclusion – Recommendations

Conclusion – Recommendations

- The use of shared devices between the RMDCN operational traffic exchange and the IPSec-based backup infrastructure created additional constraints
  - Using dedicated IPSec box should to be considered in an operational environment
- The use of IPSec devices from different vendors proved to be challenging
  - Consider using one device type or at least one device brand for an operational deployment
- “manual” re-routing is time-consuming and prone to mistakes
  - The traffic re-routing has to be fast, automatic and reliable. Only dynamic routing processes can ensure this in an operational environment

Conclusion – Next Steps

- **If agreed during ROC-14**
  - ECMWF will create a test environment for DMVPN including 6 or 7 routers
  - These routers will be used to evaluate DMVPN Phase 3 locally
  - If successful, Q4-2008 3 or 4 routers will be sent to volunteers sites to try DMVPN over the Internet
  - DMVPN will then be used to create the IPSEC VPN solution to backup the RMDCN
  - Q1-2009 results of these tests.
  - If successful, consider recommendation of Cisco Routers using DMVPN for the backup of the RMDCN
  - Otherwise, market survey to find the correct solution
- **Choice of the agreed perform to be done during ROC-15 (spring 2009)**

Thanks to

- Belgium: Rudi Swennen
- China: Lang Hongliang and Sun Haiyan
- Germany Inge Essid and Ilona Glaser
- Turkey: Ahmet Ertürk, Serdar Kocaoglu and Cemal Oktar
- ECMWF: Sebastien Barbereau

Questions?