

Carrier Ethernet Global Interconnect



Service Provider Interoperability Event Pilot Phase

presented at



Paris, February 2010

With participation of leading European service providers



Media Partner



Test Equipment Provider



EDITOR'S NOTE



Carsten Rossenhövel
Managing Director
EANTC AG

Does anybody remember the world before the Internet, when IP packet networks were islands — some larger, some smaller? A number of Internet Service Providers (ISPs) already offered great IP services back then, but accessing their networks from anywhere in the world was a major headache and usually

came with very limited service guarantees. At one point in the early 1990s, however, the situation radically improved when standards like BGP matured and ISPs were able to create dependable, large-scale peering — either at Commercial Internet Exchanges (CIX) or through bilateral connections.

In 2010, Carrier Ethernet service providers worldwide are exactly at the same point: Having set up individual service offerings, they are now looking to increase their global and regional footprint by interconnecting with other carriers.

This project aims at the core of the solution: Improving *real* Carrier Ethernet service interconnection (not just based on Ethernet over SDH/SONET circuits) by validating the end-to-end functionality, high availability, service level management and scale against a common standard.

So far, service providers found that outsourcing segments of their end-to-end Ethernet services was rather an individual effort — time consuming and expensive to start with and to maintain.

Now that the Metro Ethernet Forum (MEF) has completed the long-awaited Carrier Ethernet External Network-Network Interface (ENNI), service providers can use a well-defined specification for interconnection moving forward.

In November 2009, EANTC and ancotel invited European service providers to test their readiness to interconnect Carrier Ethernet services through ENNI. Six carriers signed up for the initial pilot testing phase. After a month-long preparation, testing culminated in a one-week session where the participants carried out detailed functional and resilience tests of all interconnections at ancotel's telehouse in Frankfurt/Main, Germany.

The results confirmed a couple of things:

- Functional Ethernet service interconnection works well today — independent of the size or geographic coverage of the carrier, Ethernet services of the participating service providers were well understood and managed.
- Interconnection was straightforward only in a few cases initially: ENNI bears a few configuration pitfalls — basic in nature but sometimes difficult to track down. Standardized procedures to validate the first time interconnect and individual circuit provisioning can be helpful to avoid expensive troubleshooting.
- Point-to-point services were supported across all participants; once the market for multipoint

services may develop in the future, support will likely grow.

- Proper selection of the Ethernet peering switches helps to ensure full support of all ENNI options supported. All equipment involved (purposefully unnamed here) was MEF 9 and MEF 14 certified; this was not a guarantee for the providers per se to pass ENNI testing easily. EANTC is editing the MEF's upcoming ENNI Abstract Test Suite to ensure that equipment passing ENNI conformance testing will actually help to reduce lab efforts for service providers.
- End-to-end high availability and performance guarantees across ENNI require thorough testing to ensure that they will work as expected.

The results of the pilot test session were promising; there is more testing yet to come with regards to service level guarantees, performance monitoring or special features such as hairpin switching. Light Reading will publish the next phase test report in time for Light Reading's *Ethernet Europe* in London mid April. Given Carrier Ethernet's fast progress, we look forward to see even more European and international service providers participating in the next stage.

We hope that you will find this short test report insightful. Any comments and suggestions are very welcome.

A WORD FROM ANCOTEL



Michael Böhlert
VP Marketing
ancotel

How to manage the massively growing Carrier Ethernet Services business in the short term and how to develop efficient, secure, robust and manageable services satisfying these new opportunities in the perspective of mid-term business development? Technical development and fast progressing standardisation

dominate the discussion of Carrier Ethernet interconnection today. Service providers, at the same time, are frequently facing the challenges of interconnecting larger and more complex multi-provider networks.

The industry strongly demands scalable and resilient Carrier Ethernet interconnections and vendors have mastered to show technical compliance on a wide scale of services in the past. Now is the right time to initiate a more practical approach on the service provider level in such multi-network environments. With EANTC, we jointly designed a test program within a realistic and pre-existent environment under genuine conditions. In ancotel's carrier hotel in Frankfurt, Germany, currently more than 370 carriers out of 59 countries are hosted and interconnected by more than thirty thousand physical and virtual cross connects, an ideal playground for this test program.

We expected the "Advanced global interconnect test program" to show us how to efficiently do both, map Ethernet services to those of other service

providers and to interconnect with many partners, while preparing for the upcoming global coverage and the immense scale of services that need to be exchanged.

By demonstrating the current state of Carrier Ethernet interconnection deployment by service providers, we have gathered a vendor-independent feedback on the practical implementation throughout this initial test campaign. The participating service providers gained valuable experience enabling them to optimize their business strategies very closely aligned with the latest market needs.

PARTICIPANTS

Service Provider	Services Tested
BSO Network Solutions, France	Ethernet Virtual Leased Line (EVLL) Lan to Lan
Level 3	Ethernet Virtual Private Line (EVPL)
Neterra, Bulgaria	
P & T Luxembourg	Ethernet Virtual Private Line
Tinet, Italy	Ethernet Extension Ethernet Private Line
Ucomline, Ukraine	Ethernet Private Line Service

INTRODUCTION

The Metro Ethernet Forum (MEF) has just approved the External Network-to-Network Interface (ENNI) specification which provides a consistent and straightforward standard for Ethernet services spanning multiple provider networks.

The ENNI specification defines the data plane of the interface between two providers, defines the service segments and serves as a common basis for negotiation of service properties between carriers. The specification introduces the concept of Operator Virtual Connection (OVC), a building block for the service segment maintained by a single operator. Connected to each other, multiple OVCs form an Ethernet Virtual Circuit (EVC) to carry an end-to-end customer service across multiple networks. From the customer point of view, the EVC will appear as a single service independent of the fact that multiple carriers collaborate to provide it.

At the ENNI, an S-VLAN tag is used to unambiguously identify a specific circuit and allow the providers to correctly associate each exchanged frame with one of the OVCs within their network. This mechanism can be used for bilateral peering over an Ethernet link between two service providers. ENNI also enables Carrier Ethernet interconnect exchanges, such as the one hosted by ancotel for this pilot test, to efficiently and flexibly construct services across multiple providers. Exchanges drastically reduce the number of individual links and arrangements required to build an interconnection infrastructure.

For the purpose of this test, we constructed a prototype of a central interconnection facility, based on a single provider-grade Carrier Ethernet switch.

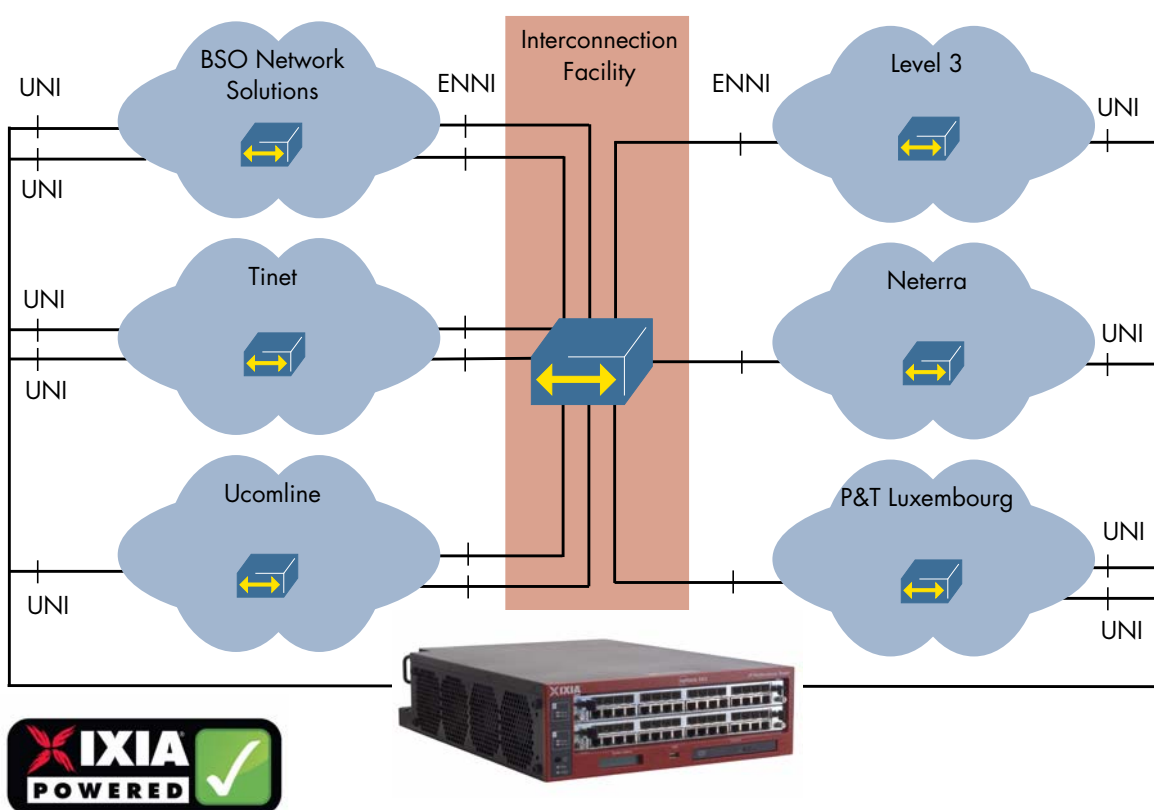


Figure Physical Test Topology (all links Gigabit Ethernet)

TEST SETUP

This section describes the test setup and configuration in detail. Test results are covered in the following section on test results starting on page 5.

The participating service providers were locally connected at ancotel's telehouse in Frankfurt, Germany. Each of the providers made at least one UNI (User-Network Interface) port and one ENNI (Enterprise Network Network Interface) port available for testing. Some of the participants provided two UNIs in order to diversify testing possibilities.

Three of the carriers provided two physical ENNI ports to test link protection using a link aggregation group. We kept the resilient configuration active throughout all the tests.

Each of the providers provisioned their services in accordance with the test and configuration plan. We defined configurations suitable to maximize the number of tested features and minimize the number of required re-configurations.

The ENNI links of all providers were connected to the Carrier Ethernet interconnection facility. The UNI links were connected to an Ixia XM2 with LM1000XMV16 modules (16 ports of 1 Gigabit Ethernet) running IxNetwork to emulate enterprise customers by sending and receiving appropriately formatted Carrier Ethernet traffic. Finally, one port of the Ixia XM2 was connected to the interconnection facility directly to allow for pre-testing of individual circuits (UNI to ENNI); this port emulated an ENNI port of a peer service provider using proper ENNI encapsulation.

Before performing the actual tests, we manually pretested each of the providers' OVC configuration manually. We ensured Ethernet frame forwarding from UNI to ENNI, correct packet encapsulation and use of proper S-VLAN and CE-VLAN IDs.

Test Combinations

During our test preparation, the participating providers were queried regarding the technical features of their network and whether they are able to support different kinds of Ethernet services. In this test phase we considered only point-to-point services. Across all MEF standards, services differ depending on the way customer Ethernet frames are mapped at the UNI and encapsulated as they enter the Metro Ethernet Network (MEN). Service segments in the MEN are called Operator Virtual Circuits (OVCs). At the UNI, customer frames are mapped either according to the Customer Edge VLAN identification (CE-VLAN IDs) or the 12-bit IDs of the Customer VLAN (C-VLAN) tag. Alternatively, all traffic arriving at a port can be mapped to a single service regardless of any tags.

Depending on the service, the CE-VLAN ID may or may not be preserved during the transport. Obviously when multiple CE-VLAN IDs are mapped to the same OVC on a single port, they must be preserved in order to identify the VLAN correctly.

We tested four types of customer tag mappings:

- **Mapping A:** All CE-VLAN IDs are mapped to the OVC and preserved, resulting in a transparent VPN service for the customer.
- **Mapping B:** A single CE-VLAN ID is mapped to an OVC and preserved. This type of service may be useful for connecting multiple branch offices to headquarters over the same port at the central location to distinguish branch office traffic.
- **Mapping C:** A single CE-VLAN ID is mapped and not preserved. The CE-VLAN ID is only used for local traffic selection at the UNI.
- **Mapping D:** Multiple CE-VLAN IDs are mapped and preserved.

At the ENNI, another tag is added called the Service VLAN (S-VLAN) tag. Its identification (the S-VLAN ID) is used to identify the Operator Virtual Circuit across the ENNI. This way, VLAN IDs used by the customer are left unaffected and identical IDs used by different customers can coexist on the same interface.

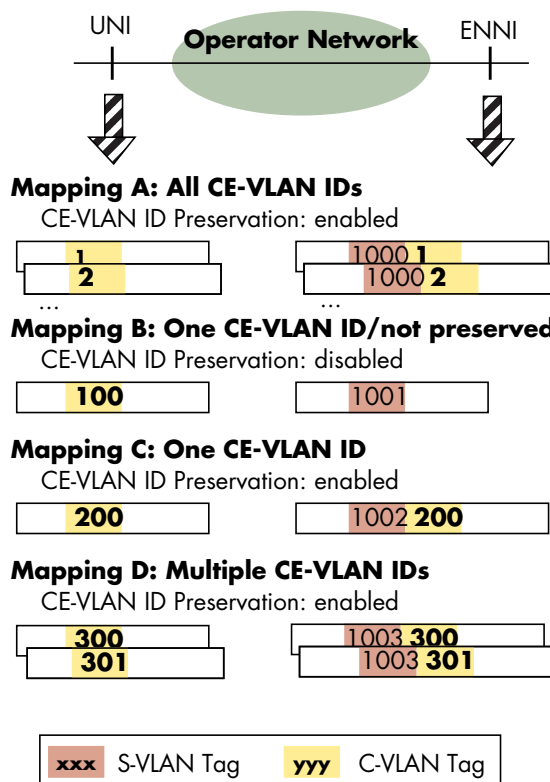


Figure 2: Types of Mapping (Frame Format)

End-to-End Scenarios Evaluated

For each Ethernet service scenario, we selected all possible pairwise combinations of providers supporting a particular type of customer traffic mapping.

In each test, we sent bidirectional end-to-end unicast traffic between the UNIs of the two service providers involved. We used frame formats matching the respective service requirements as described above. Given that the tests were functional and did not exceed the service guarantees, we expected traffic forwarding without any packet loss. Nevertheless we measured packet loss and latency using the Ixia

XM2.

Since all services were point-to-point Ethernet connection, no Ethernet address (MAC) learning was performed before sending the traffic. The traffic consisted of Ethernet frames without any higher layer protocols — this was specifically selected to ensure packets were forwarded independent of any Internet Protocol (IP) addresses or IP-based filtering.

We used an IMIX-style frame size breakdown to emulate a real world scenario. Every group of 12 consecutive frames consisted of seven 64-byte frames, four 570-byte frames and one 1518-byte sized frame. In total we generated a bandwidth of 10 Mbit/s each time. (Performance objectives and service levels will be tested in the next phase.)

Protection Setup

In this pilot test phase, we focused on the standard link protection mechanism specified by ENNI, which is based on link aggregation.

As a side note, there are several mechanisms to protect Carrier Ethernet services against node failure which EANTC has tested in previous vendor interoperability events at Carrier Ethernet World Congress and MPLS World Congress since 2007, for example based on CFM, G.8032 and MPLS. These will be covered in future phases of this service provider test program.

The ENNI specification requires the use of a special type of link aggregation across ENNI (requirements R1, R2):

- There must be exactly two physical Ethernet links
- Exactly one link must be active, while the other link must be in standby mode and only be used upon failure of the primary link
- Load balancing must be disabled.
- Link Aggregation Control Protocol (LACP) must be in use.

These options are supported in IEEE 802.3-2005 ("802.3ad"), now formally moved to IEEE 802.1AX-2008 — however, they are not the default options.

BSO Network Solutions, Tinet and Ucomline participated in this test. We connected their ENNIs to the central interconnection facility using two Gigabit Ethernet links combined into a link aggregation group. Link protection was implemented separately at each ENNI. This setup allowed for flexible link protection where supported; end-to-end connections could be set up with all six providers independent of link protection support or ENNI connection across a conventional single Ethernet link.

During the test preparation, we found that support for Ethernet link aggregation was implemented on all the equipment involved, but vendors and providers were still working on installing the particular options of link aggregation required by ENNI. We decided to setup the link protection in a similar, albeit not strictly ENNI-conformant way. As a result, the protection was realized with a conventional link aggregation including load balancing

across the two parallel links. Upon failure of one of the links, link aggregation was expected to switch all traffic onto the remaining link.




Figure 3: ENNI Protection Setup

Protection Test Scenarios

The protection functionality on the ENNI was tested separately for each provider supporting this feature in order to isolate the network under test given the complexity of such a setup. Without the second provider present in this setup, the analyzer was connected to the interconnection facility directly and emulated the ENNI of the other provider.

We sent bidirectional Ethernet traffic between the UNI and the ENNI, including the S-VLAN tag. In the service configurations without CE-VLAN ID preservation, we excluded the C-VLAN tag on the ENNI as described on page 4.

In this test, we used the same traffic mix as above. The test duration was increased to five minutes.

Each test run started with both links connected and bidirectional traffic initiated. After approximately two minutes, we manually interrupted the first link in the link aggregation group by pulling the link. Afterwards, we continued to send traffic until the end of the test with only one link available. We continuously recorded the per-second statistics of the frames received by the Ixia XM2 on both the UNI and ENNI ports to evaluate the impact on the service caused by the interruption.

TEST RESULTS: BASIC INTERCONNECT

We performed all basic interconnect test cases with the participating service providers with relative success. In the preparation phase, we encountered several configuration issues which we were able to resolve with the providers during the test period (Jan 8–19, 2010) quickly.

Initially we observed S-VLAN ID configuration issues.

Equipment used by some of the participating service providers maintained a global S-VLAN ID table where the IDs required in our test plan were already in use on other ports of the remote equipment. As a result, we had to coordinate S-VLAN IDs across service providers, ensuring that a given ID for a service would be unique not only across the ENNIs in our test but unique in all the provider networks involved.

Due to lack of support for proper 802.1ad encapsulation (Ethertype 0x88a8) at some providers' equipment and inability of the interconnect switch to flexibly compensate for this problem per port, we resorted to a non-conforming Q-in-Q encapsulation at all ENNI ports using Ethertype 0x8100.

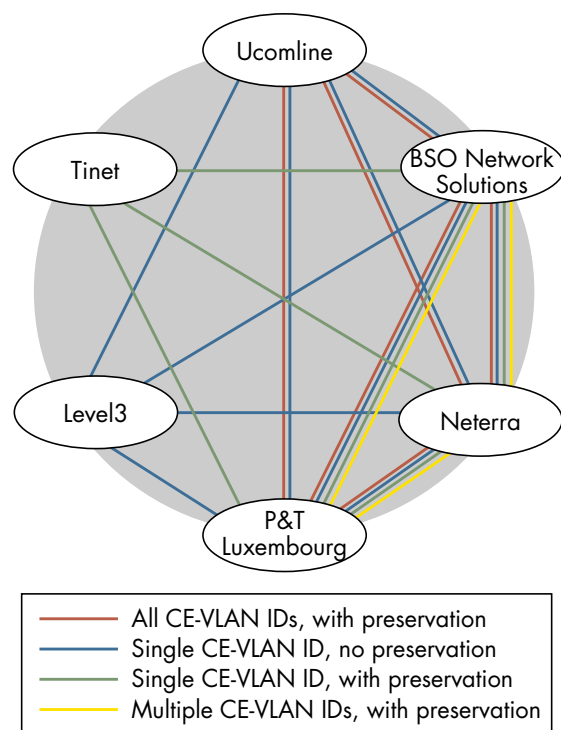


Figure 4: Basic Interconnect Combinations

The figure above depicts all test combinations and connection types tested.

During the tests, we identified two problems described below, which in each case could be attributed to a specific network. In a real service, these are however unlikely to cause a significant impact on the service.

In one case, we encountered an issue with tagged 64-byte frames transmitted at the UNI and over OVC without CE-VLAN ID preservation. The C-VLAN tag was stripped from the frames at the ingress, which caused the frame size to be reduced to 60 bytes, below the minimal size allowed in Ethernet.

Although according to the provider, the untagged frames were never transmitted over physical links directly, the equipment enforced the minimal frame size by adding four zero bytes to the payload. This problem did not occur in scenarios where CE-VLAN ID preservation was enforced. The provider was able to successfully investigate and correct this issue.

In another case, we detected a minimal frame loss (less than 20 lost frames per million transmitted) and a relatively high latency of 30-40 ms. According to the provider, this can be attributed to the need to use a leased line to forward the traffic from ancotel's location to the provider's central offices where the Carrier Ethernet switch equipment was located.

A frame loss ratio better than 10^{-5} was maintained by all providers except the one affected by the packet loss. But even in the latter case, the frame loss ratio was still below 10^{-4} .

TEST RESULTS: ENNI PROTECTION

We were able to successfully perform all protection tests with the three providers that offered support for this feature.

Due to the inability of the interconnection facility to support the an LAG-based protection with an active and a standby link as required per ENNI specification, we had to resort to using a conventional link aggregation. Nevertheless, existing mechanisms of link aggregation were able to redirect all traffic to the remaining link after interruption of one of the links.

In all cases, we observed end-to-end service loss for a failover period of three seconds or less following the link failure. Packet forwarding was fully restored afterwards. In one case, the restoration initially took 30 seconds and could be fixed by changing the LACP timeout setting to "fast" therefore lowering the LACPDU transmission interval from 30 to 1 second.

SUMMARY

All service providers involved were able to overcome the initial problems within the test interval of only two weeks. We were able to demonstrate perfect functional interoperability in all supported Carrier Ethernet point-to-point service combinations in a basic Ethernet service interconnection scenario.

Lack of support for standby links in link aggregation groups in some equipment prevented us from using a protection mechanism fully compliant with the ENNI specification. Nevertheless, the conventional link aggregation proved to be a reliable workaround in our interoperability scenario.

The upcoming next phases of the ancotel / EANTC Global Interconnect test campaign will focus differentiated services and service levels, multipoint services and hairpin switching as well as fault and performance monitoring across end-to-end services with European and worldwide service providers.

PARTICIPANTS

About ancotel

ancotel GmbH, founded in 1999 and headquartered in Frankfurt am Main, operates the largest and most significant telecommunications and data hub in continental Europe with more than 360 service providers co-located and interconnected in its neutral carrier hotel. As an independent company, ancotel is a leading provider of managed services in the area of virtualized carrier transport and interconnection infrastructure for legacy and NGN networks, including TDM, VoIP, IP/MPLS and Carrier Ethernet.

About EANTC



The European Advanced Networking Test Center (EANTC) offers independent telecom network test services for manufacturers, service providers and enterprise customers. Business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP, Carrier Ethernet, Triple Play, and IP applications.

About BSO Network Solutions

Founded in 2004, BSO Network Solutions is the leading Next Generation Operator providing Network services, Hosting and Integration. With its rapid growth, BSO Network Solutions is already present in more than 12 countries (France, UK, The Netherlands, Germany, Belgium, Italy, Spain, Switzerland, USA, Brazil, Canada and Hong Kong) and has offices in Paris, London and Hong Kong. BSO Network Solutions administers and operates its own Nx10Gbit/s network which has been crafted to be compatible with future 100Gbit/s interfaces. Thanks to a range of very high value-added services and to comprehensive facilities management and validation offerings, BSO Network Solutions provides its customers with guidance and assistance from the most upstream phases of their projects, in terms of advice and expertise, to the most downstream, including the daily administration, upgrading and adaptation of their infrastructure. All its services are covered by extremely robust Service Level Agreements (SLA) tailor-made to match each problem addressed.

www.bsonetwork.com

About Level 3

At the core of an increasingly networked world, Level 3 Communications (NASDAQ: LVT) is a premier international provider of fiber-based communications services. We are a proven carrier for enterprise, wholesale, government and content customers, who rely on Level 3 to deliver advanced solutions with an industry-leading combination of efficiency and quality. We own and operate one of the world's most scalable, end-to-end networks to connect our customers with a robust portfolio of metro and long-haul services, including transport, data, Internet, content delivery and voice. Learn why Level 3 is the trusted provider for the largest users of bandwidth:

www.level3.com

About Neterra

Neterra is a leading Bulgarian telecommunication operator and provider of high quality terrestrial and satellite communication solutions. Besides the sale of large Internet capacities to corporate clients and ISPs it constructed the communication networks of branches of world financial institutions and Bulgarian banks.

At present Neterra imports the major part of the international Internet traffic into Bulgaria. It offers unique solutions for monitoring of the networks of its clients, a platform for internet television and broadcasts television and radio signals in the country.

Neterra owns and operates one of the most highly technological communication centers in Bulgaria – Sofia Teleport. The services are reliable, the maintenance – timely, and the work standards - universally recognized.

<http://www.neterra.net/en>

About P&TLuxembourg

P&TLuxembourg is the number one telecom operator in Luxembourg, with a turnover of 349 million Euros in 2008. As a comprehensive operator and owner of its fixed and mobile infrastructure, P&TLuxembourg builds and operates with TERALINK one of the most innovative European IP networks offering customers data and telecommunications solutions and services by providing connections to Europe's foremost telecommunication and carrier hotels.

www.teralink.biz / www.pt.lu

About Tinet

Tinet, formerly the carrier arm of Tiscali Group, is the only global carrier exclusively committed to the IP/MPLS wholesale market.

With network presence and customers in EMEA, Americas and APAC, Tinet provides global IP Transit and Ethernet connectivity to Carriers, Service and Content Providers worldwide, within 7 working days. The carrier guarantees customers proactive management of SLAs and protection from DDoS attacks.

Established in 2002, Tinet's unique business model, based on focus and simplicity, assures the delivery of the highest standard of service. Tinet has grown to become one of the top 10 global IPv4 backbones and the number one IPv6 network worldwide.

www.tinet.net

About Ucomline CJSC – Member of VEGA TELECOM Group

VEGA TELECOM is the largest alternative fixed operator in Ukraine providing full scale of telecom services in fixed retail and wholesale markets, serving more than one million subscribers in retail telephony and data services. UCOMLINE is a wholesale division of VEGA TELECOM Group.

VEGA TELECOM owns and operates a nationwide fiber-optic network, DWDM backbone and robust metro networks as well as a highly developed and modern MPLS network.

www.vegatele.com



EANTC AG
European Advanced
Networking Test Center

Einsteinufer 17
10587 Berlin, Germany
Tel: +49 30 3180595-0
info@eantc.de
<http://www.eantc.com>



ancotel GmbH

Kleyerstr. 90
60326 Frankfurt, Germany
Tel: +49 69 750013-200
<http://www.ancotel.de>



Light Reading
United Business Media

11 West 19th Street
New York, NY 10011,
USA
Tel: +1. 212.600.3000
<http://www.lightreading.com/>



Upperside Conferences

54 rue du Faubourg Saint
Antoine
75012 Paris - France
Tel: +33 1 53 46 63 80
info@upperside.fr
<http://www.upperside.fr>

This report is copyright © 2010 EANTC AG. While every reasonable effort has been made to ensure accuracy and completeness of this publication, the authors assume no responsibility for the use of any information contained herein.

All brand names and logos mentioned here are registered trademarks of their respective companies in the United States and other countries.

20100215 v0.5