INF-3996

MASTER'S THESIS IN TELEMEDICINE AND E-HEALTH

The Scandinavian Health Network: Connecting the Scandinavian countries' health networks

Anders Baardsgaard

SUBMITTED:
June 1st 2007

Faculty of Science
Department of Computer Science
University of Tromsø

Preface

After building regional and national health networks full time for eight years the time felt right in 2005 to do something else, and I chose to embark on a two-year master study in telemedicine. This thesis marks the end of these studies, and two years leave of absence from Norsk Helsenett AS. As in all aspects of life, this background has had positive and negative consequences for the task at hand: to say something sensible about how to build a Scandinavian health network. I try to focus on the positive.

I have realized for some time that some of the technical solutions selected in Nordnorsk Helsenett, many of which later to be adopted in Norsk Helsenett, have limitations. Although well suited for employment in a limited, regional setting, they have scaling issues that make integration with other networks difficult – be it network mergers or connecting to other, autonomous networks. I have tried to describe some of these shortcomings, and suggest some alternative paths.

One of the most positive experiences during this work is the realization that there are other national health network organizations and infrastructures out there, beyond a basic knowledge that such networks exist and what their names are. There is a potential for common benefit in increased contact and exchange of experiences and viewpoints, that may lead to improved solutions on a national level, as well as providing a fundament for building cross-border health networks. In this respect, the health sector networks may also have some lessons to learn from network organizations in the academic sector.

Acknowledgement

I would like to thank my tutors professor Gunnar Hartvigsen and assoc. professor Johan Gustav Bellika for essential guidance and advice in the process of writing this thesis. Discussions with Dr. Gunnar Klein and assoc. professor Stig Kjær Andersen of Aalborg University were also helpful.

There are many who deserve to be thanked in Norsk Helsenett AS. First of all, CTO Vidar Eriksen and CEO Mette Vestli for granting me leave to undertake the master study under the conditions that I requested, and for liberal working conditions during the past two years. Friends and colleagues at Norsk Helsenett AS' operations and support centre in Tromsø (all of you), but in particular Roger Jørgensen and Ørjan Wasmuth for good discussions on central issues. And thanks to Morten Celius, Hans Petter Krogsrud and Hans Jørgen Varfjell at the main office in Trondheim for sharing information with me.

Several colleagues in Carelink, MedCom and affiliates have been very patient and helpful in providing me with information and feedback, and conducting measurements: Sten Haglund, Claus Duedal Pedersen and Ib Lucht.

Fellow Telemedicine students deserve to be thanked; Gudleif, Harald, Odd-Arne, Sindre, Terje, Ieva, Kristoffer, Frode, Andrei, Kristian and Håvard. Both for the parties and the odd relevant discussion, but most of all for the invaluable feeling of being in the same boat.

And thanks to my wife and kids for their patience. I really look forward to being able to return to the regular schedule, with more generous time frames for common activities like weekend excursions, assembling trampolines, bicycling and other important stuff!

Summary

The Scandinavian countries are all among the minority of the world's nations that have established a national health network infrastructure. These networks provide a secured communication environment for health professionals, mainly in hospitals and local health services/GPs. Despite serving similar requirements, these national infrastructures differ in several aspects: organizational extension, choice of carrier technology, IP addressing and NAT strategy, DNS implementation and QoS support.

Initiatives to interconnect these national health networks in Scandinavia date several years back, and today the Danish health network has a number of cross-border connections, including VPN tunnels to the Swedish and Norwegian health networks. However, limitations in the connecting technology, and scalability issues related to technological architecture, pose limitations on the cross-border service provisioning.

This thesis investigates the differences in technical infrastructure, discusses consequences of these differences, and suggests modifications to harmonize the technology among the national health networks involved. Some of these suggestions may even have merit within national borders, as they can be seen to address scaling issues in intranetworking in general.

Contents

List of Figures

List of Tables

1 Introduction

1.1 Background

Exchange of digital data has as long a history in the health sector, as in the rest of society. Although is may seem that health institutions in general have been comparatively slow in embracing digital technologies in medical applications, the trend towards computerization is now accelerating. Data in the health sector is increasingly collected, stored, refined, evaluated and exchanged in digital form, displacing paper-based systems.

But what is the geographical scale of the operation? Do we need to exchange health data across national borders? What is the rationale for investigating technological options in interconnecting national health networks?

Cross-border telemedicine projects have been going on for many years, illustrating the need to exchange digital medical data and services between institutions in different countries. Some of these activties are briefly described in section 2.6. Also, interconnection of national health networks already has some merit, although the history is not so long (see section 2.4.6).

This trend towards increased cross-border communication in the health sector is also promoted by the EU commission. Their communication document to the Commission and The European Parliament titled "e-Health – making healthcare better for European citizens: An action plan for a European e-health Area" [EUComm04] states that:

> Increased networking, exchange of experiences and data, and benchmarking, is also necessary at the European level in the health sector. Drivers for this include the need for improvements in efficiency, and the increased mobility of patients and health professionals under an emerging internal market in services. The situation requires the integration of clinical, organisational, and economic information across health care facilities, so as to facilitate virtual enterprises at the level of jurisdictions and beyond.

Additionally, the study of technology for trans-border interconnection of health networks may even provide feedback to the task of creating smaller networks, on a regional or national level. Scaling issues are relevant and worth considering in most cases when planning and building technical infrastructures.

1.2 Health networks in context

Joining of larger data networks is no new or rare undertaking. Any company merging process is bound to face questions concerning how to integrate the participating parties' ICT infrastructure and service delivery into a common, shared unit. With the restructuring of the public health sector that took place in

Norway in 2002, and the one that is made effective in Denmark from January 1. 2007, "company mergers" may be seen to occur within the health sector as well. The intended result of such a merging process is a single organization, under a single management, where the ICT is an integrated part of this unity.

1.2.1 Sector networks

Connecting organizations into sector networks is a related activity, where the intention is not a full integration into a new enterprise, but rather to share a limited selection of data and services between several organizations. These organizations continue to be independent even after the genesis of the sector network. The purpose of the sector network is to create an arena for limited exchange of ICT data and services between the connected companies. Banking and finance have a long tradition with such cooperation, fuelled by the fundamental requirement to transfer monetary units between them. The Norwegian Oil sector network initiative SOIL established in the mid-1990s is another example. Reliability, security and dependability are frequently listed as motivation for not using the Internet for traffic exchange in sector networks [SOIL07, Arvidsson06].

Regional, national, and forthcoming international, health networks also fall into the sector network category. But while organizations exchange data via a few well-known applications in many other sectors, the health sector displays a requirement to exchange data via an increasing number of applications, many of which are not predefined [Pedersen05].

1.2.2 Describing health networks

Any sector network is bound to be shaped by the sector that it is created to support. Obviously, the professional applications differ from sector to sector, as does the legal and organizational framework for the connected organizations. [Nohr05] has outlined some of the legal aspects that form a basis for this framework with regard to the health sector, discussing issues like privacy, confidentiality, responsibility and licensure. These issues in turn are reflected in national laws and regulations, which have an impact on both what kind of organizations that may be connected, how they should be organized, and the prerequisites for exchange of (sensitive) data over the health network. The Norwegian Sector Norm [SHdir06] may serve as an example of a very concrete set of requirement for health network connected organizations.

The Norwegian Department for Health and Social Affairs suggests partitioning the IT based information and communication systems into three levels [HoD96]: applications, information infrastructure and technical infrastructure.

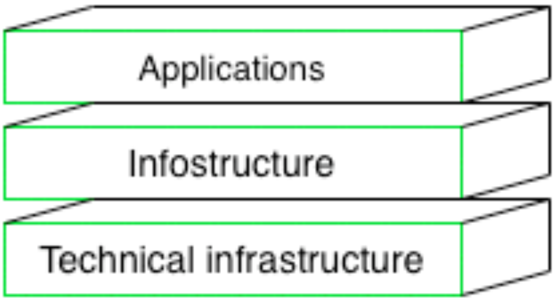

Figure 1 Health network levels

The prime motivation for establishing a health network is to provide a secured arena for use of *applications* to exchange electronic health services [NHN07a]. The first generation of health networks was limited to services based on asynchronous message exchange, like laboratory results, discharge notes, medical prescriptions etc. This service category still forms an essential part of the service exchange in health networks [MedCom06a]. More mature health networks additionally enable user access to interactive services, frequently web based access to data stored in a database, although many protocols and formats may be used. Example services include PACS/RIS and EHR/EPR. A third service category, provided by modern health networks, is real time services with strict service quality (QoS) requirements. Video conferencing and IP telephony are the most common QoS dependent services.

While electronic health services for medical professionals are exchanged between health organizations, some directory and dictionary services are provided by the health network organizations themselves, and provide data of a less sensitive character. The underlying contents of this *infostructure* is frequently extracted from data originating in the health organizations, as is the case with DNS [RFC1034, RFC1035]. Other directories may be actively populated by the health organizations, like national catalogues of health personnel and resources [NHN07b, Carelink07a, Sundhed05]. These catalogues may be designed to serve purposes like

- identifying service requestors and providers
- authenticating users of interactive services
- enabling encryption and electronic signatures (PKI)

Important features in the realization of the *technical infrastructure* in a health network are IP address management and NAT strategy. These issues are related, as the most common cause for use of NAT is that private IP addresses are used internally in a network. These issues are important because they influence the robustness of the communication sessions transmitted across the network, as well as the service diversity that the network is capable of handling, and the ability to interconnect to other networks with a minimum amount of configuration and operational hassles.

1.3 Problem definition

The main problem of this thesis is: *Can a Scandinavian Health Network infrastructure be built from the infrastructure components of the national health networks*?

The information publicly available on the technical implementation details in the national health networks is lacking and fragmented. Particularly the Norwegian Norsk Helsenett has little material published. A survey in this area, covering technological platforms and services provided, will provide a platform for further comparison, alignment and cooperation

The current level of interoperability between the established national health networks should be described, and the potential for future interconnections and their technological requirements should be discussed. Of particular interest are current and future real-time capabilities (QoS).

The feasibility, utility and desirability of a Scandinavian Health Network should be addressed, as well as the requirements that the network's technical infrastructure should fulfill.

1.4 Main results

For the first time, a description of the basic technological elements of the infrastructure in Norsk Helsenett is collected and made publicly available, together with similar features in the other Scandinavian national health networks, and a comparison is made. Also, performance measurements are executed between these networks. Both of these achievements rest heavily on the author's unique background in building health networks, and access to human and technical resources in the health network organizations.

The status of the infrastructures is analyzed and discussed in light of the basic functionalities that each of its components ideally should provide. Several deficiencies of the technology currently used are pointed out and discussed. The corrective measures proposed should be carefully considered, even for purely internal functionality improvements, before embarking on an effort to build a common, full-functional Scandinavian health network infrastructure.

It should be noted that a large proportion of the sources cited are of a non-academic character. In fact, very few academic sources exist on the issues discussed in this thesis. Also, a number of sources are cited as "personal communication", where interviews or email was used to obtain additional information on undocumented issues.

1.5 Outline

The structure of this thesis is as follows:

Chapter 2 presents some infrastructure components of health networks. Vital aspects of the national health networks in the Scandinavian countries are described and compared. The status of national health network infrastructure in some additional countries, as well as some cross-border e-health applications, is briefly described.

Chapter 3 discusses network performance measurement methodology, and applies this to the trans-border health network infrastructure in Scandinavia, to describe a measurement scenario. Unfortunately, re-allocation of external resources mandated changes in this scenario, with unforeseen consequences.

Chapter 4 presents motivational discussions on the infrastructure components covered in chapter 2, and suggests requirements in each area. Novel ideas include use of IPv6 centrally assigned local addresses and introduction of a new public top-level domain .health.

Chapter 5 presents the results of the network performance measurement conducted between the national health networks, and offers some interpretations.

Chapter 6 offers a discussion of the main structural, technological and motivational issues involved in creating a Scandinavian health network.

Chapter 7 presents some concluding remarks,

2 Health network infrastructure

2.1 IP addressing and NAT

The original Internet Protocol specification from 1981 [RFC791] briefly describes the 32 bits IP address format, together with the initial separation into A, B and C class networks. At the time there was hardly anyone who could anticipate the proliferation that the technology was to achieve, and the dominating role it was to obtain in the market. But from the mid-1990s it has been clear that the original IP address space would be insufficient for the projected growth in use of Internet technology, under the requirement that each host should have a unique permanent address taken from the specified range. Some worry was also expressed with regard to the growth in size of routing tables in the Internet's core routers, but the main consideration was directed towards the rapid and accelerating consumption of IP address space.

The administration of the IP address space is currently delegated from The Internet Assigned Numbers Authority (www.IANA.org) to five regional Internet registries (RIRs), each handling IP address allocation requests for a part of the globe. In Europe, Réseaux IP Européens (www.RIPE.org) plays this role. The others are APNIC, covering the Asia/Pacific region, ARIN, covering North America, LACNIC, covering Latin America/Caribbean and AfriNIC, covering Africa. Each RIR handles requests for allocation of public IP address ranges according to regional policies, e.g. [RIPE-388, RIPE-405].

2.1.1 Network address translation (NAT)

In 1994 a proposal was put forward to reuse IP addresses in order to reduce the growth rate in IP address allocations. The original document named "The IP Network Address Translator (NAT)" [RFC1631] was superseded by a more mature proposal in 2001 [RFC3022]. It is to be regarded in conjunction with [RFC1918] which reserves three address ranges that will never be allocated to any organization for official Internet use, and thus may be safely (re)used for private/NAT purposes. Together, these two documents form the basis of a strategy to allow the current Internet technology to continue to be used in an ever increasing market. The non-private IP addresses are called public addresses.

NAT was designed to let a potentially large number of hosts share a smaller number of network level addresses when communicating externally, i.e. with hosts that provide services on the Internet. This is accomplished by replacing the 32-bit private IP address of the internal host with an IP address from a smaller pool of available public addresses. This address translation needs to be executed by a NAT device that is situated on the perimeter of the internal network. It must be in a position to intercept internally originating IP datagrams destined for the Internet, modify the source IP address, and forward the datagram on a valid path towards its destination.

In order for data traffic to be bi-directional, the NAT device is also required to intercept datagrams in the reverse direction, changing the public IP destination address in the incoming datagram to the private IP address of the original datagram. To accomplish this, it is necessary for the NAT device to maintain a mapping between the private IP addresses of the internal hosts that are currently communicating with Internet hosts, and the (usually) dynamically allocated public IP addresses that their Internet based communication partners see.

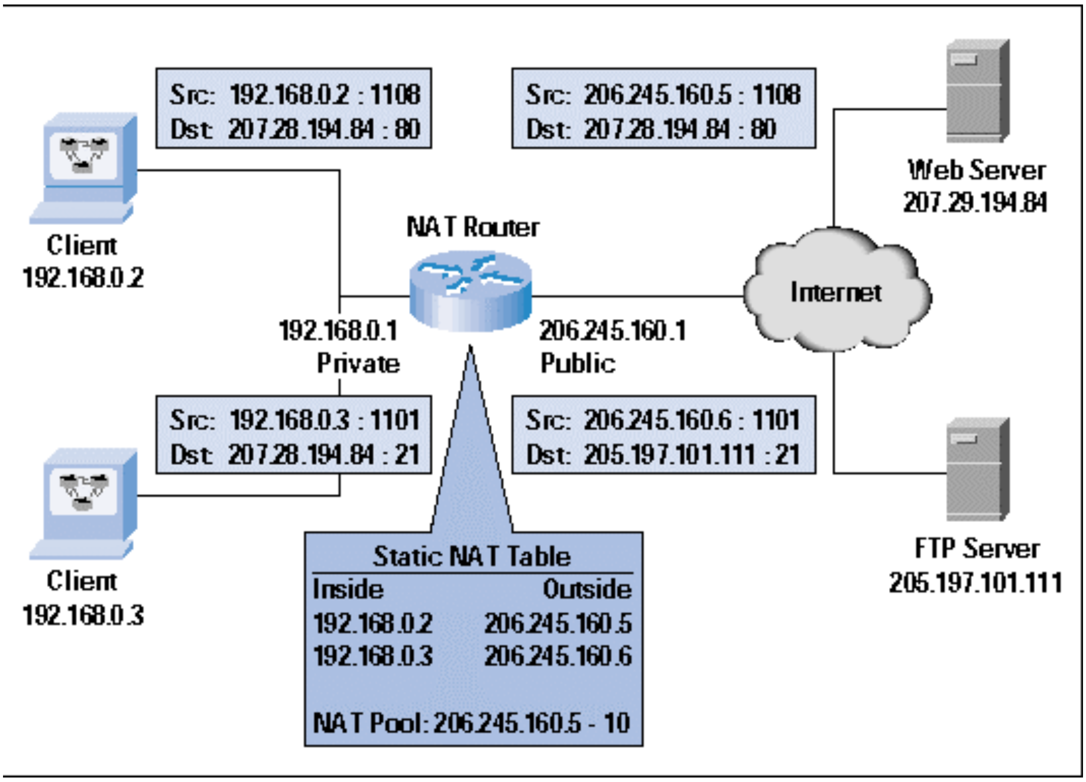

Static NAT Table

Inside	Outside
192.168.0.2	206.245.160.5
192.168.0.3	206.245.160.6

NAT Pool: 206.245.160.5 - 10

Figure 2 Dynamic NAT [Phifer00]

One of the refinements described in [RFC3022] is the ability for several internal hosts to simultaneously share a single public address in their communication with Internet hosts. In addition to the IP address modifications prescribed by NAT, this requires similar modifications to higher-level protocol parameters like TCP and UDP port numbers, and was called Network Address Port Translation, NAPT.

[RFC2663] discusses NAT in the context of address realms, which is defined as "a network domain in which the network addresses are uniquely assigned to entities such that datagrams can be routed to them". This uniqueness property is the locator functionality of addresses as described in [RFC2101]. NAT devices basically attempt to provide a transparent routing functionality to end hosts trying to communicate from disparate address realms. Transparent routing differs from traditional routing in that IP address contents of the IP header is modified.

2.1.2 Negative implications of NAT

Although helpful in managing the growing demand for IP addresses, it was clear from the outset that NAT employment would also pose some problems. An important consequence of the transparent routing functionality is that it weakens the end-to-end principle, originally introduced in [Saltzer84] but also fundamental to Internet technology [RFC1958]. This principle is generally recognized to bundle desirable consequences as protection of innovation, reliability and robustness [RFC3724]. The main problem with NAT in this respect is that it increases the amount of state in the network above the minimum level that the network needs to perform its services, such as routes, QoS guarantees etc. "This state must be self-healing (…) The volume of this state must be minimized, and the loss of the state must not result in more than a temporary denial of service given that connectivity exists." ([RFC1958] p. 3) The implication is that state ideally should be maintained only in the endpoints. Such state will only be lost when the endpoint itself fails. [RFC2993] details some facets of keeping state in the network:

- difficulties in routing around problems
- state in the core will tend to grow with the network, potentially creating severe choke points due to capacity problems in the individual elements
- if security is included in the state then the possible trust models that the network can support become restricted.

[RFC2775] gives two additional accounts of end-to-endness: performance and address transparency. Degradation of the second of these properties is the primary concern: "IPv4 addresses can no longer be assumed to be either globally unique or invariant, and any protocol or application design that assume these properties will fail unpredictably."

[RFC3027] goes more in details regarding the problems that some common Internet protocols encounter when they exchange traffic across a NAT device. As stated in the document the list is not comprehensive. Neither is the discussion of solutions and workarounds regarding each protocol exhaustive, as the commonly used option to tunnel X11 applications through ssh is not mentioned. Instead, the X Windowing system is sorted under the heading "Protocols that cannot work with NAT enroute", together with IPsec and IKE, Kerberos 4 and 5, and rsh/rlogin. Protocols that can be made to work with an application level gateway (ALG) or proxy solution include FTP, RSVP, DNS, SMTP, SIP, RealAudio, H.323 and SNMP. It is also noted that in general, peer-to-peer applications are even more likely to break with NAT enroute than client-server applications.

As a practical aside it may be noted that the inherent property of NAT to cause address collisions between hosts in different address realms has already been reported to give operational problems in the Norwegian health sector. [Lother05] contains an account of a situation where message exchanged between Norsk Helsenett and Tromsø municipality was impeded by an IP address conflict. Potential remedies of a more general nature to such address collisions exist, like twice NAT [RFC2663] and Realm Specific IP [RFC3102], but they introduce additional complexity in the network configuration.

NAT is usually employed at the network perimeter, frequently in conjunction with a firewall. Many vendors implement NAT and the security features of a firewall in the same hardware, which contribute to a blurring of the two functionalities. Furthermore, NAT is sometimes presented as a security feature in itself because the opposite scenario, where public internal addresses are used, might help potential intruders if known externally. [RFC 2775] points out that this argument is false, since it is trivial to hide addresses by suitable access control lists: “A system with a hidden address is just as private as a system with a private address.”

Finally it should be noted that firewalls share some of the undesirable features with NAT devices, in that they introduce additional state in the network, and tend to separate a network into smaller domains with restricted inter-domain communication. But in contrast to NAT, firewall functionality may be perceived to have an intrinsic value, as it is widely accepted as a security feature. “It should however be noted that there will always be administrative boundaries, firewalls and intranets, because of the need for security and the implementation of policies. NAT is seen as a significant complication on these boundaries.” [RFC2956]

2.2 DNS

The basic functionality of the Domain Name Service (DNS) is to provide a directory lookup from host names to IP addresses, a process called *resolving*. In the first ARPANET host names were resolved locally on each computer by means of a hosts file. In principle this file contained the name and IP address of every host connected to the network. But the task of keeping the hosts file updated on all hosts soon proved to be increasingly complex, as the connection rate tended towards exponential growth.

Following a meeting in 1982 to discuss addressing issues in computer mail, it was decided to introduce multi-level naming (domains) [RFC805]:

```
The conclusion in this area was that the current
"user@host" mailbox identifier should be extended to
"user@host.domain" where "domain" could be a hierarchy of
domains.
```

This process concluded in the specification of the Domain Name System (DNS) in 1987 [RFC1034, RFC1035]. An illustration of the importance of DNS addresses in contrast to use of IP addresses is the following quote from [RFC1900]:

```
To make renumbering more feasible, the IAB strongly
recommends that all designs and implementations should
minimise the cases in which IP addresses are stored in non-
volatile storage maintained by humans, such as
configuration files.  Configuration information used by
TCP/IP protocols should be expressed, whenever possible, in
```

> terms of Fully Qualified Domain Names, rather than IP addresses. Hardcoding IP addresses into applications should be deprecated. Files containing lists of name to address mappings, other than that used as part of DNS configuration, should be deprecated, and avoided wherever possible.

The modern Internet DNS is a distributed, hierarchical database, It provides lookup functions on a number of different *resource record* categories (RRs), but the most important are still the A RR for lookup of a host's IP address from its name, and the PTR RR for lookup of a host's name from its IP address.

The lookup functionality is provided to DNS clients by *name servers*, which come with a number of different per-*zone*[1] capabilities: primary, secondary, forwarding and non-authoritative [Albitz01]. Relevant for the current presentation are the primary name servers, each serving one or more zones. Sub-zones are delegated to cooperating name servers, forming the DNS hierarchy, whereas no hierarchical structure is imposed on the name servers themselves[2].

When private IP addresses are used, DNS data pertaining to these hosts are required to be contained within the enterprise [RFC1918]. But as the benefits of a functional DNS is no less on an internal enterprise network than out on the open Internet, this usually gives rise to a *split horizon* DNS for the enterprise. The main feature, and drawback, of a split horizon DNS is that the zone contents is not the same on the inside (the internal network) as on the outside (the Internet): the answer to a DNS query depends on whether it is sent from the inside or the outside.

Technically, a split horizon occurs when a name server acts as a primary server for a zone, without proper zone delegation from the parent zone. The consequence is a disconnected DNS system at the "kidnapped" zone, where the properly delegated zone is part of the Internet DNS' *global name space*, while the kidnapped zone forms (part of) an internal DNS' *internal name space*.

Another DNS characteristic is the *common root*, where all qualifying resource records are gathered in a sub-tree of the DNS name space. The conecept actually consists of two parts: a uniqueness property and a characterizing property. The *uniqueness* property states that all relevant DNS resources should share a common DNS suffix, and the *characterizing* property states that all resources that share the common suffix actually qualify as relevant. In the current context, the relevancy criterion is that the resources belong to the health sector, for some interpretation of "belong".

[1] A DNS zone has the same contents as the corresponding domain, minus any delegated (sub-)domains.

[2] A name server A may be primary for both the zones a and a.b.c, while delegating the zone a.b to another server B which in turn delegates the zone a.b.c back to A.

2.3 Quality of service

The term quality of service (QoS) generally refers to mechanisms used to provide different priorities to different streams or classes of network traffic, in order to be able to guarantee bounded values for the quality parameters bandwidth, packet loss, delay and delay variation [Sørensen04].

The Integrated Services Architecture (IntServ) [RFC1633] is based on an extension of the traditional best effort model, to support real-time service of IP. It requires the network to provide special QoS for specific user packet streams (flows) and uses a resource reservation protocol (RSVP, [RFC2205]) to enable the application to signal and reserve the desired QoS from the network. The IntServ model requires the routers to maintain flow-specific state, a strategy that has obvious scaling issues for large networks.

The Differentiated Services Architecture (DiffServ) [RFC2475] presents an alternative, where traffic is classified into a limited selection of service classes, marked and possibly shaped by boundary routers on ingress. The traffic is routed through the network by interior routers that operate according to a predefined per-hop-behaviour per service class, that describe the characteristics of the forwarding. On egress the traffic is handled by boundary routers which deliver the traffic to destination hosts.

There is also a proposal to combine the two architectures, to achieve the best from both worlds [RFC2998].

A review of QoS metrics is postponed to section 3.1.1, in connection to a description of the performance measurements conducted between the national health networks.

2.4 National health networks in Scandinavia

All three Scandinavian countries have established national health networks based on Internet technology. The following text presents an account of their extent, topology, basic technology and services.

2.4.1 Differences in health sector organization

The Scandinavian countries have a long history of cooperation, and the societal similarities are large. This also applies to the health sector, although differences exist. One such aspect, the responsibility for and ownership of the health organizations, is useful to be aware of as background knowledge to understanding the structure of the current national health networks.

At the end of the previous millennium, the counties were responsible for the specialist health care in all three countries. But in 2002 the Norwegian central government took ownership of the public hospitals, and established five wholly owned regional health enterprises that in turn own the hospitals (via health enterprises). In 2007 a similar reorganization was undertaken in Denmark, whereas Sweden has maintained the counties' role as hospital owners.

The organization of the primary health care also displays some differences. In Norway and Denmark, primary health care is a municipal responsibility, although many of the GPs are self-employed and operate under contract with the municipal authorities. Whereas in Sweden, only care of the elderly is a municipal responsibility, while the county level is responsible for other health areas.

2.4.2 Denmark

MedCom, the Danish Health Network project organisation, is currently in its fifth generation and project portfolio. It was established in 1994 [CFST07], and was made permanent with MedCom-III in 2000 [MedCom07a]. From the outset, the focus was on format and transport of EDI messages, but with MedCom-IV in 2002 interactive service capability was specified through The Internet Strategy sub-project. This was implemented in 2003 [MedCom06b]. MedCom is currently financed in a cooperative effort by both national, regional and local health authorities, as well as the national pharmaceutical association. The staff is employed at The Center for Health Telematics in Odense, and reaches a head count of approximately 15.

2.4.2.1 Infrastructure and DNS

The Internet based Sundhedsdatanet (Health Data Network) is presented as a supplement to the traditional Sundhedsdatanet, which is the original VANS network for transport of EDI messages [MedCom04]. A third component for electronic communication in the Danish health care, The Health Portal (Sundhedsportalen, www.sundhed.dk), mainly concerns communication between patients and the health care system, but also lists sub-goals directed towards health professionals [Siticom02].

The core of the Internet based Sundhedsdatanet (I-SDN) is the Sundheds-DIX (SDN). The technical realization of the network is by GRE tunnels inside IPsec VPN connections in a hub-and-spoke configuration over the Internet, i.e. having an ordinary Internet subscription with static IP addresses from an ISP is a prerequisite for connecting to SDN. The VPN tunnels terminate in equipment configured with IP addresses from the ISP, leaving the encapsulated GRE tunnel unencrypted. The GRE tunnel in turn is unpacked either on the same equipment as the VPN tunnel, or on a separate unit closer to the connected organization's internal network core [Sorth03, Cisco01].

The I-SDN network uses public, provider independent IP addresses, and has been allocated the range 192.80.240.0/20 from RIPE. These addresses are not routed in the Internet, and hence not accessible outside I-SDN [Bech03]. Each connected organization is allocated a subnet of this range, and is expected to configure NAT for traffic that is routed across I-SDN, although sub-subnets may be statically configured on devices that are used for special purposes, e.g. video conferencing [MedCom07b].

Domain names in the I-SDN network are registered under the artificial .medcom top level domain. The DNS system is detached from the Internet DNS, and connected organizations are required to forward DNS queries regarding .medcom or 240-255.in-addr.arpa resource records to the central MedCom DNS servers.

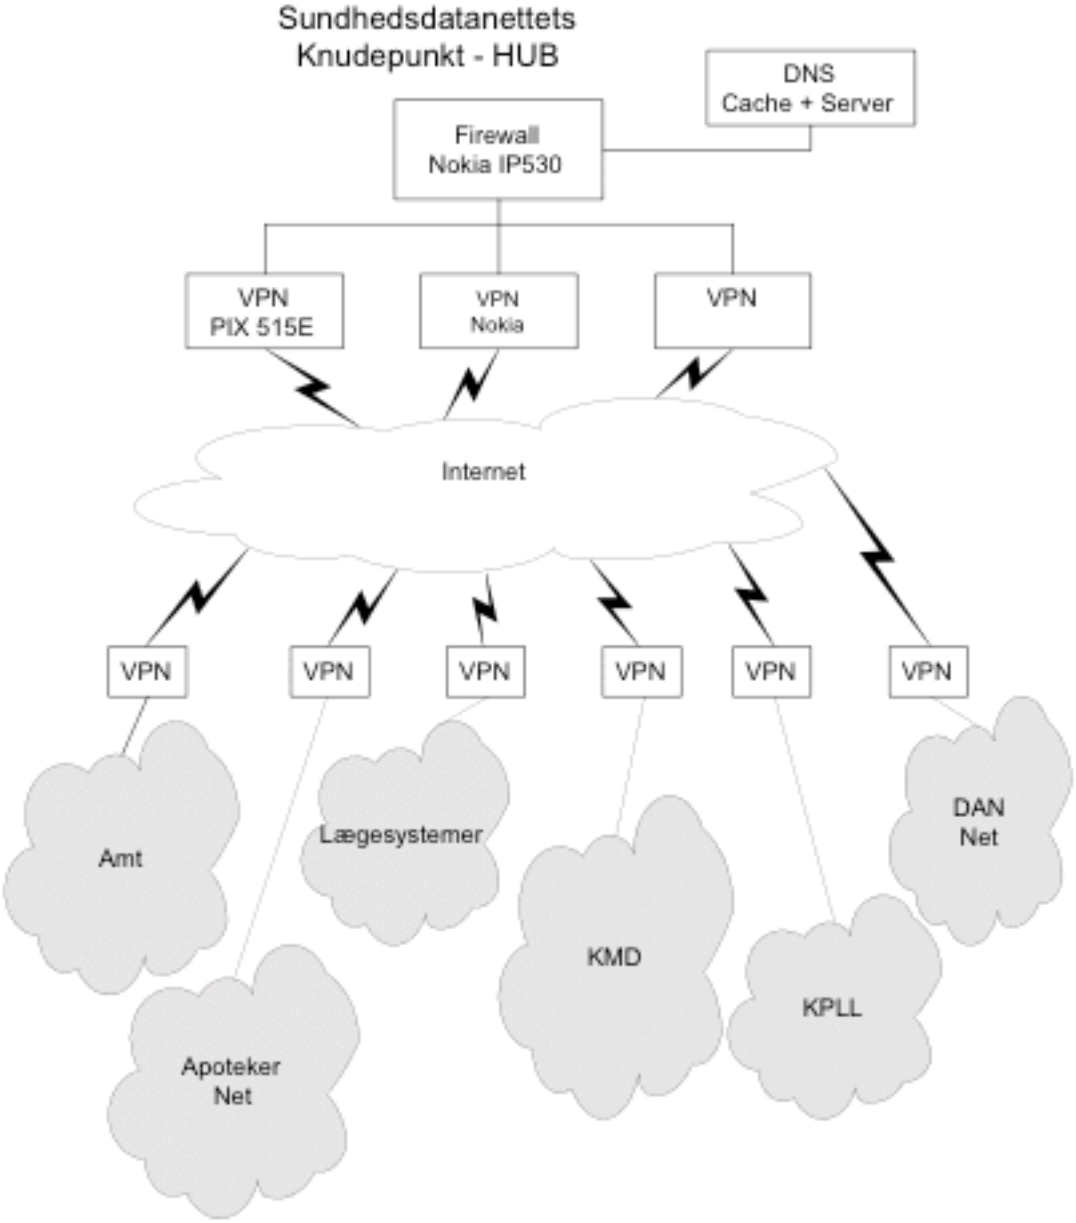

Figure 3 The Internet based Health Data Network [Bech02]

2.4.2.2 Connected organizations

According to [MedCom07c], 50 organizations have VPN connections to the I-SDN by April 2007. But this figure includes all the health regions that own the public hospitals, KMD which serves the municipalities with ICT services, and the pharmacy network. The actual coverage of Danish health organization is therefore considerably higher: in reality, all Danish hospitals, GPs, municipalities and pharmacies are connected to I-SDN, and have the capability of exchanging interactive services over the network. The actual number of traffic exchange contracts in the connection agreement system (Aftalesystemet, [Mogens07]) is in excess of 1500 by February 2007 [Pedersen07].

2.4.2.3 Operation and services

Operations and maintenance of the I-SDN is performed by UNI-C. Also included are the internal DNS service and an MCU for videoconferencing service. In addition, a large number of medical and other services are available from the connected organizations. One feature of the Aftalesystemet is that it gives a good account of available services, which organizations provide them, and from which network addresses. [MedCom07d] lists such services grouped in 14 categories, both medical and other, with an average of 10 service providers/access points in each.

2.4.3 Sweden

The Sjunet project was initiated by seven Swedish county councils in 1998 [Malmqvist04]. Since 2001 Carelink, a limited company that works with ICT in Swedish health care, has been responsible for Sjunet, in close co-operation with all the county councils and representatives for the private care providers and local authorities. Carelink was founded by local, county and national health authorities, in cooperation with private care providers and the pharmacies. Carelink's activity is financed mainly through service fees from member organizations, although one third of the 2006 budget was direct government funding. The organization has 12 employees in Stockholm.

2.4.3.1 Infrastructure and DNS

The Sjunet infrastructure is currently in its third generation. After the initial implementation as a VPN network provided by Telia, and a prolongation period (Sjunet 2, 2000-2), the current infrastructure was acquired from Song Networks in 2003 [Carelink04]. Sjunet is implemented as an IP-VPN overlay to Song's core network, consisting of gigabit Ethernet links arranged in overlapping VLANs, with OSPF based IP routing between them, for redundancy and increased availability [Carelink05a]. Sjunet members connect to this network with individual access capacity as required.

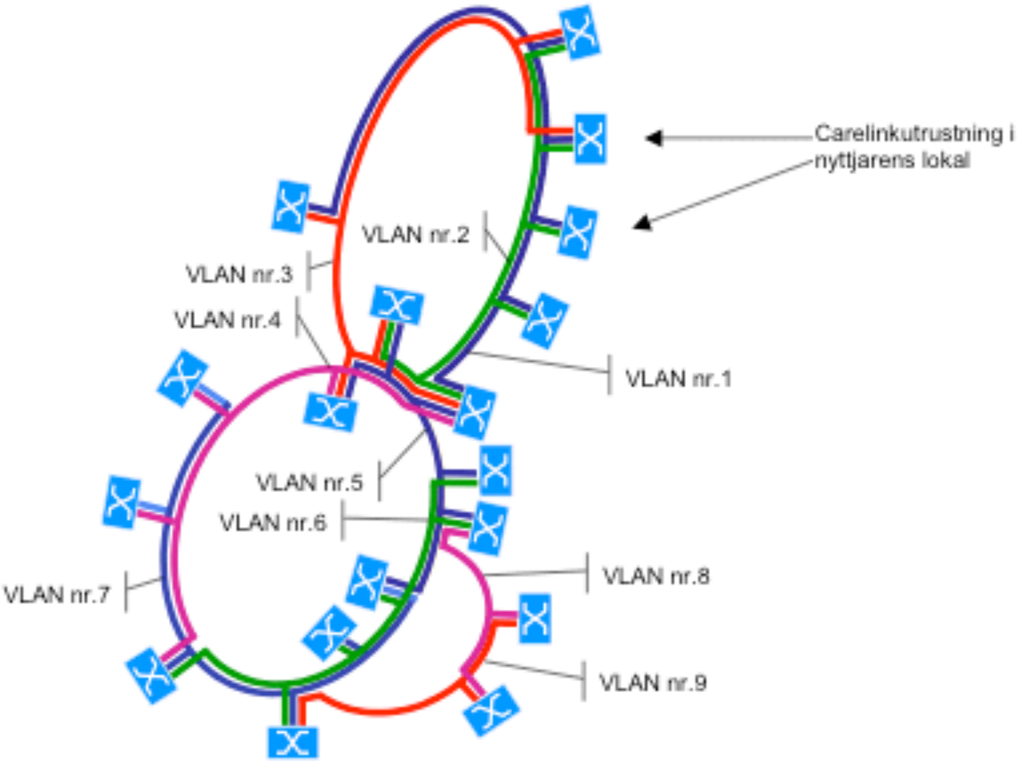

Figure 4 Sjunet physical accesses to Song core VLAN topology [Carelink05a]

Only public IP addresses are routed over Sjunet. Carelink members in lack of such addresses are allocated a limited IP range by Carelink from one of the blocks 82.136.128.0/19 and 213.189.96.0/19, which Carelink has acquired through its LIR membership with RIPE NCC. If this IP range, or the range of public IP addresses that the connected organization has acquired from other sources, is insufficient to the number of hosts and internal network topology, then NAT is required for network traffic bound for Sjunet [Haglund07].

Sjunet uses the domain sjunet.org as root in its DNS tree. This domain is also available on the Internet, with the same contents as internally in Sjunet. All Sjunet connected organizations are required to construct their own sjunet.org domain by removing the .se suffix from their Internet domain name, and prefixing the remains to .sjunet.org. They are further recommended to use the same resource names internally and within Sjunet [Carelink05a].

2.4.3.2 Connected organizations

All the Swedish counties (Landsting) are members of Carelink, and connected to Sjunet. The public hospitals are in turn connected to the counties' internal networks, giving a connection rate of 100%. As noted above, the Swedish primary health care is also a responsibility for the counties. The implication is that all GPs, except a small number of doctors in the company health service (Företagshälsovården), are connected to Sjunet. Furthermore, 42 municipalities and 7 private care providers are connected [Carelink07d]. Finally, the national pharmacy monopoly's network is connected to Sjunet.

2.4.3.3 Operation and services

Services available through Sjunet are operated by a range of service providers, under contract with Carelink [Carelink07b]. Purists consider Sjunet to consist of only the basic network access and IP routing service, which is delivered to Carelink member institutions by Song Networks. The DNS service for sjunet.org is operated from The Academic Hospital in Uppsala. Another important Carelink service is HSA, a national directory of organizations, personnel and roles in the health sector. Technically integrated in HSA is SITHS, a PKI service based on personal ID cards with integrated electronic certificates, issued to health personnel. Other Carelink services include a video conferencing MCU, as well as electronic support services provided by private enterprises and central public institutions to the health sector [Carelink07c].

2.4.4 Norway

In the second half of the 1990s the five health regions established separate regional health networks, with no inter-regional coordination. Most aspects of these creations differed – including technology, service ambitions and organisation – but each of them eventually converged into a basis of Internet technology. In 2003 the National Health Network project was established by the Directorate for Health and Social Affairs (SHdir), to establish a central infrastructure that would connect these five regional health networks. This six-network conglomerate, including equipment, personnel and contracts, was transferred to a new limited company Norsk Helsenett AS in 2004. The company is owned in equal shares by the five regional health enterprises, and currently employs 40 people at the headquarter in Trondheim, operations and support center in Tromsø and branch office in Oslo. Norsk Helsenett AS is financed mainly through service fees from connected organizations.

2.4.4.1 Infrastructure

The conglomerate legacy of Norsk Helsenett is readily visible in the network's topology. Most hospitals and other organizational units in the specialist care service still have their connection to Norsk Helsenett via the old regional health network structure. A contract was signed with Telenor for an IP-VPN service in October 2006 and the central infrastructure has been moved to their Nordic Connect platform. The intention is to phase out the regional networks as well, and have the hospitals connect directly to the same carrier network. Although a transition process has commenced, the old regional health network infrastructures still dominate. "Having the specialist care institutions' health network connections converge to a unified platform, and reduce the number of transport network service providers and carrier technologies, is high on our agenda for developing Norsk Helsenett as a tool for the health sector. The operational simplifications in leaving the conglomerate of partly self-managed and partly outsourced infrastructures are obvious. Additional benefits include optional redundant network connections to the institutions, bandwidth offerings scalable to gigabit capacity, and end-to-end QoS in the network", says Vidar Eriksen, CTO of Norsk Helsenett AS.

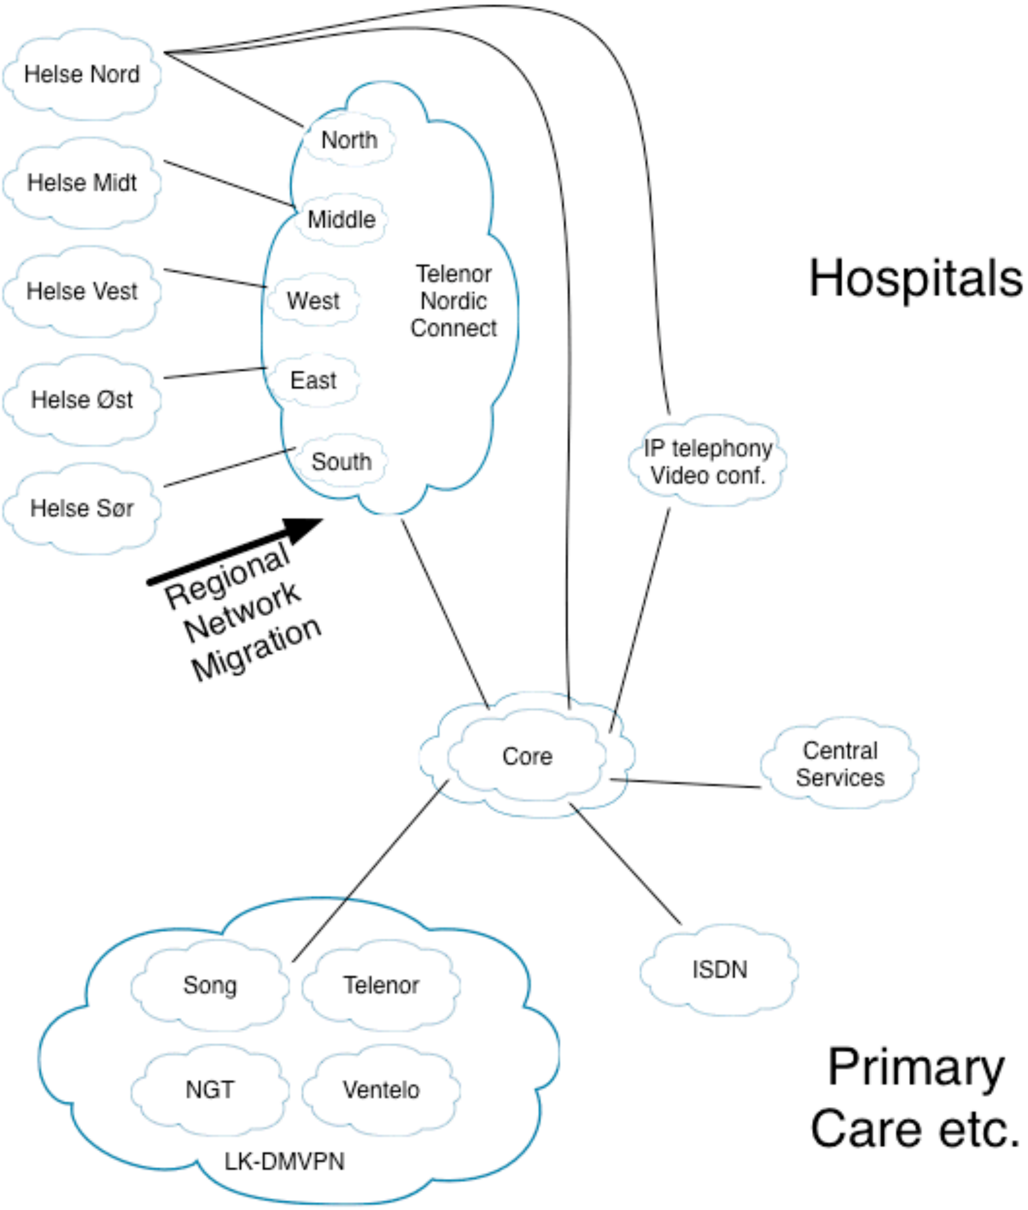

Figure 5 Norsk Helsenett network components

The choice of IP addressing strategy also dates back to the National Health Network project, where representatives from all five regional health networks participated. It was decided to emphasize end-to-end principles, and avoid NAT as far as possible. A national IP address plan for the health enterprises and GPs was worked out, allowing both private and public addresses to be routed in the network as long as the public address ranges used were officially registered with the institution by RIPE or other regional internet registry [SHdir03]. NAT was discouraged but not prohibited, as the workload involved in renumbering a large number of hosts in the most "unfortunate" organizations could be very large. Besides, municipal care was not included in the plan, as the ICT functions including IP address regime of such units would be managed by the ICT department in their respective municipal administrations.

In 2006 Norsk Helsenett AS entered into a LIR membership with RIPE NCC and has acquired the IP range 91.186.64.0/19. The main use of this block is to overcome some of the address conflict issues arising when municipalities connect to Norsk Helsenett. Two scenarios are remedied;

i) NAT of internal IP addresses before they enter the health network to avoid address collisions and enable internal services to be reachable from the outside, if desirable; and

ii) allocating public addresses to central services, avoiding the possibility of internal clients trying to access them as local services.

2.4.4.2 DNS

The DNS in Norsk Helsenett uses Internet DNS as a basis, but masks out zones for organizations connected to Norsk Helsenett, and replaces them with zones containing health network internal resource records. This internal DNS system is operated in a manner similar to the Internet DNS, with organizations maintaining zone data (resource records) on their own DNS servers, and central slave servers to download zones and act as forwarders.

In the design document for the National Health Network project, a zone listing mechanism for DNS servers was included. This was intended to remove the requirement of manual zone transfer configuration from the central slave servers, but this feature was abandoned by Norsk Helsenett. The figure below is included primarily to illustrate the complexity of the scheme, while the actual procedure described is less important.

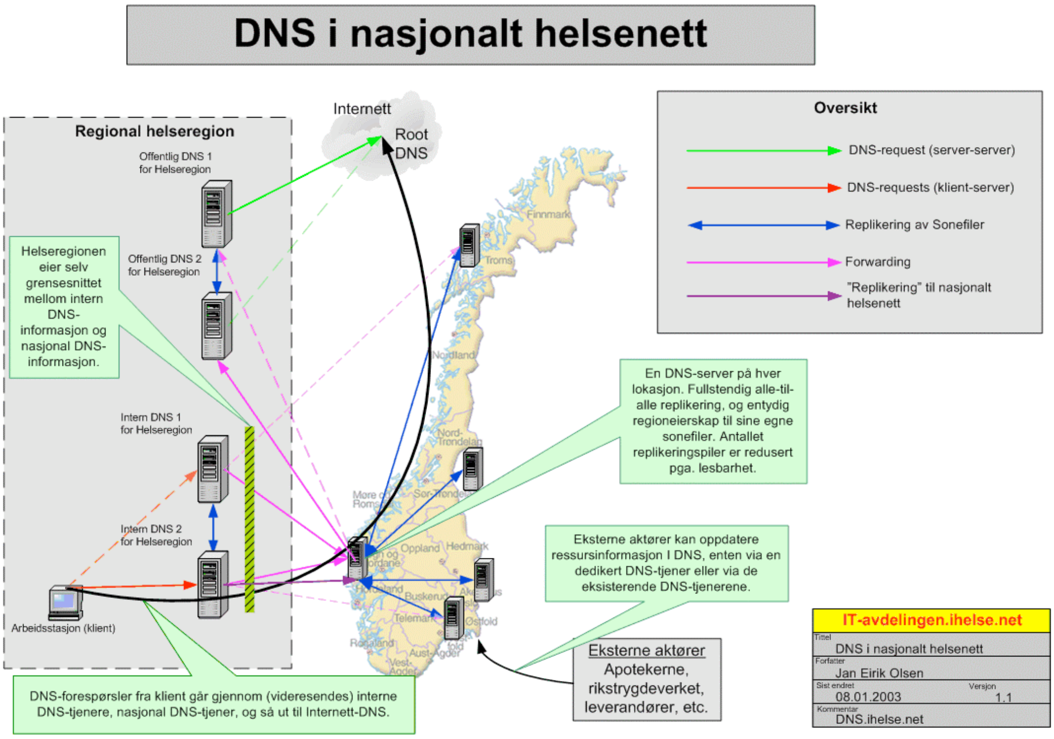

Helseregionen eier selv grensesnittet mellom intern DNS-informasjon og nasjonal DNS-informasjon.

En DNS-server på hver lokasjon. Fullstendig alle-til-alle replikering, og entydig regioneierskap til sine egne sonefiler. Antallet replikeringspiler er redusert pga. lesbarhet.

Eksterne aktører kan oppdatere ressursinformasjon I DNS, enten via en dedikert DNS-tjener eller via de eksisterende DNS-tjenerene.

Eksterne aktører
Apotekerne,
rikstrygdeverket,
leverandører, etc.

DNS-forespørsler fra klient går gjennom (videresendes) interne DNS-tjenere, nasjonal DNS-tjener, og så ut til Internett-DNS.

IT-avdelingen.ihelse.net
DNS i nasjonalt helsenett
Jan Eirik Olsen
08.01.2003
1.1
DNS.ihelse.net

Figure 6 DNS as planned in the National Health Network project [SHdir03]

2.4.4.3 Connected organizations

From the inception of Norsk Helsenett AS, the organization was granted an exclusive right to deliver network infrastructure to the regional health enterprises by the Department of Health. This right implies both a requirement on regional health enterprises and their subsidiaries to acquire their communication infrastructure from Norsk Helsenett AS, and an obligation for Norsk Helsenett AS to deliver the infrastructure and communication services requested by the health enterprises. All Norwegian hospitals are connected to Norsk Helsenett.

For GPs and specialists in private practice the connection rate is 60%, while 22% of the municipal health services are connected [Krogsrud07]. All the pharmacies communicate via the pharmacy network, which in turn is connected to Norsk Helsenett.

2.4.4.4 Operation and services

Operation of Norsk Helsenett is the responsibility of the operations and support centre in Tromsø. For what remains of the regional health networks, management is outsourced for four of them, while the IP infrastructure of the old Nordnorsk Helsenett is operated in-house. Also insourced is the part of network connecting GPs, municipalities etc. In addition, central services and connections to external service providers are managed from Tromsø.

In addition to the basic IP transport and DNS services, Norsk Helsenett offers a number of services, as well as connections to external service providers. The Address Register (Adresseregisteret, previously HER) is a directory service much like Carelink's HSA, with integration of a currently outsourced PKI functionality planned. There is a centralized service for exchange of EDI messages, a coordination system for transport of patients and an MCU based video conferencing service. Optional services include access to Internet world wide web and email, either content filtered or via terminal server. External service providers are made available for remote management, ASP, POS terminals, as well as various contents providers, both private enterprises and public registers.

2.4.5 Comparison of the health networks

The national health networks in Scandinavia have been implemented with very little contact between the technical personnel specifying and implementing the infrastructure in each country. This lack of technical coordination has led to a number of differences in the organization and operation of the networks. Some of these differences may have consequences for network interconnections.

While the Swedish and Danish health networks mostly connect intranets that in turn connect health institutions, the Norwegian health network extends to the individual hospitals' and GPs' network connection point. One consequence of this is that the number of connection points to Norsk Helsenett runs two orders of magnitude higher than Sjunet and I-SDN.

The service regimes of the networks differ. Norsk Helsenett poses restrictions on members' (direct) connections to other networks, and the service portfolio includes filtered Internet services to make this less painful. This "single datacomm service provider" ambition is not paralleled in Denmark and Sweden.

The Swedish Landstings have responsibility for large societal areas besides health matters, like schools and public transport, and Sjunet also carries network traffic on behalf of these non-health institutions [Carelinke05b]. In Denmark and Norway, the national health networks carry only health related data.

In Denmark the use of the connection agreement system imposes a strict framework on the traffic exchange, and even integrates data encryption into the service, while in Sweden and Norway the formal framework concerning traffic exchange is less rigid. The Swedish and Norwegian health networks present requirements of a more general nature on connected organizations' use of ICT, through the Guidelines for Security [Carelink03] and Sector Norm [SHdir06] documents respectively.

The transport networks use different basic technology: Internet based VPNs in Denmark, MPLS-VPN in Sweden and several technologies in Norway, but mainly an IP-VPN service for the health enterprises (hospitals). This gives the following QoS status for the networks:

- Denmark: not possible
- Sweden: not implemented
- Norway: implemented as part of the TN-NC contract, but currently not used

The IP addressing strategy differs. The Danish and Swedish health networks both route only public IP addresses that are not routed on the Internet. The implication is that NAT is mandatory for organizations where the number of externally communicating hosts exceeds the number of IP addresses in the allocated range. In Norway, the emphasis has been on end-to-end connectivity, and there exists an IP address plan that includes the IANA allocated private IP address ranges [RFC1918].

The choice of DNS implementation differs. In Denmark and Sweden the "forward" resource records are registered under a common root; .medcom and .sjunet.org respectively, while the DNS implemented in Norway lacks this common root property. However, this has been recognized as a shortcoming, and an effort is being implemented to partly remedy the situation by registering copies of essential resource names under the nhn.no domain [Hætta07].

2.4.6 Nordic international health network

In 2005 a report to the Nordic Council of Ministers [NCM05] from a working group with representatives from the five Nordic countries described the creation of an interconnection between the national health networks in Denmark, Sweden and Norway. The network was established in 2004-5, and is based on the Danish I-SDN technology, where VPN tunnels over the Internet are used to connect Internet access points in Sjunet and Norsk Helsenett to the Sundheds-

DIX, which forms the core of the Danish health network. One view of this first Scandinavian health network is that it connects the Swedish and Norwegian health networks to the Danish, in line with Danish health regions and other health organizations in Denmark. It should be noted that this interconnection of the national health networks has a bias, in that it does not enable bi-lateral exchange of data traffic between health organizations in Sweden and Norway without the cooperation of the Danish health network operator.

Another Nordic/Baltic initiative uses the same technology to connect single hospitals in Estonia and Lithuania to the Danish Sundheds-DIX [Baltic07]. The Baltic eHealth project has established radiology and ultrasound pilots over this network.

HDN.eu is a third international initiative based on the Danish Sundheds-DIX technology [HDN07]. While it is currently in the process of seeking partners in 10 other European countries in order to apply for EU funding, it seeks cooperation with both national and regional health networks "to develop a system that will ensure **interoperability** on a transborder, transnational level between the countries' networks".

2.5 Other national health networks

There is an abundance of information on health networks – a Google search on the term turns up more than a million hits. But very few of these are actually about health network infrastructure in accordance with the present agenda. To qualify as a national health (data) network, a project, initiative or organization should:

- have a national perspective, with the aim to connect health service providers in all inhabited areas of the country
- have a sector-wide coverage, at least targeting hospitals and GPs
- coordinate technical infrastructure for secured communication between the connected organizations
- support telemedicine applications, like EDI exchange, directories etc

Besides the Scandinavian countries, the health network in **England** and **Scotland** are the only national health networks in Europe found to qualify under the criteria listed above. In England, a migration from NHSnet and the previous service provider (Cable & Wireless) to N3, an MPLS-VPN infrastructure managed by BT, was completed as recently as April 1, 2007. Among the changes made was a shift in IP strategy towards use of private IP addresses only, and employ renumbering or NAT to conform to this policy [Divaharan05]. Another notable change was the introduction of QoS capabilities [Fitchett05]. The DNS used is of the split horizon variety, and Internet services are provided to connected organizations via a combined proxy and NAT Internet gateway [Glenholmes07]. By March 2007, there were 18,689 connections to N3 in England, covering 98% of the GPs. In Scotland there were 2032 connections [Conn#5].

On the other side of the Globe, **New Zealand** commenced piloting of health networking activities as early as the mid-1990s. The country officially launched the Health Intranet in 1999, which was renamed to the Health Network in 2005 [HISAC06]. Health institutions etc connect via one of two virtual provider networks operating and linked nationally, to gain access to the National Health Index, databases, administrative systems and other services. According to one of these providers, more than 70% of New Zealand's GPs use their RSD service[3] to make referrals to hospitals and specialists [HealthLink02].

Australia has established the HealthConnect programme, where the Australian Government is investing $128 million over four years to implement "a major platform for reforming health care delivery in Australia". Although according to its FAQ pages on the Internet, "HealthConnect is neither a data repository nor a network" [HCFAQ07], it includes the Broadband for Health program, "to support the uptake of broadband services in General Practices and Aboriginal Community Controlled Health Services (ACCHS) nationwide". Supported services include Medicare Australia, clinical messaging and HealthConnect electronic health records initiatives [AG-DHA06].

2.5.1 Countries without national health networks

After spending some time investigating which countries have established a national health network infrastructure according to the criteria listed above, it seems clear that most countries have not. However, the task of presenting an up-to-date account of even the countries that come closest is difficult, given the abundance of information available on "health networks" and the fact that most of this information does not concern infrastructure at all. Therefore, the following is not written under an ambition to a present an exhaustive survey of health network status in all countries.

The WHO's Global Observatory for e-health has recently reported on its first global survey on e-health [WHO06], in which infrastructure was one of the seven thematic areas addressed. Unfortunately, the presentation makes no contribution to assessing the status of health network infrastructure development in each country. E.g. there is no mention of a national health network infrastructure in the entry on Denmark, although the nationwide web-based health portal is described in positive terms. Also, the enclosed list of WHO member states and associate members shows that only 112 of 194 members responded.

At an HDN.eu pre-project meeting in Copenhagen in January 2007 [HDN07], some representatives presented the health network status of their countries. One conclusion is that the countries **Finland**, **Lithuania**, **Serbia**, **Spain** and **Germany** do not have national health networks operational. **Iceland**, which together with the Scandinavian countries and Finland is included in the previously mentioned report on e-health to the Nordic Council of Ministers from 2005 [NCM05], has a strategy to have a health care network fully operational by the end of 2006. The current status of this ambition is not confirmed.

[3] Referrals, Patient Status Reports and Discharge Summaries

In **USA**, a report from The President's Information Technology Advisory Committee aims at modernizing the nation's health care through use of ICT [PITAC04], The report proposes a framework for a 21st century health care information infrastructure, in which one of the four focus points is a "secure, private, interoperable, electronic health information exchange". The report lists 12 recommendations, including a suggestion to promote securely encrypted, inexpensive Internet connections instead of "expensive, largely obsolete communication links". Also, regulatory impediments to e-mail communication between willing patients and their caregivers should be removed.

Neighbouring **Canada** established a National Collaborating Centre for Infrastructure, Info-Structure and New Tools Development in 2004 [CDNNCC07], but the centre's achievements after its establishment are not well publicized. The general impression is that infrastructure is not high on the list when the term "health network" is on the agenda. E.g. according to the link "Government Health Partners" on the Public Health Agency of Canada's web site "The Canadian Health Network is a national, bi-lingual Internet-based health information service funded by the Public Health Agency of Canada" [PHAoC07].

2.6 Cross-border e-health

There exist a number of telemedicine experiments, projects and activities that involve participants in two or more countries. Some are rather spectacular, like the Lindbergh operation mentioned in section 4.2.3, while others are more modest in their public display. Some seem to be aimed at "proving" that cross-border medical cooperation and exchange of medical information is possible, while others promote cases from a sensibility perspective, where such exchange has a potential for economic savings and service delivery improvements compared to current status. It seems natural to focus on the latter category here.

One of the most experienced participants on the international e-health arena is MedCom, with a history of such activities that dates back to 1996 [MedCom07f]. The list of international e-health initiatives that MedCom is involved in includes the current Nordic health network, the eBaltic project and the HDN.eu initiative mentioned above. The infrastructure established by the eBaltic project is used for teleradiology between a Danish hospital and two hospitals in the Baltics, and for teleultrasound between a Swedish and a Norwegian hospital [Wanscher07].

Historically, teleradiology has been one of the driving forces in implementing health networks within national borders. On the international arena there are several accounts of hospitals purchasing teleradiology services from TMC in Barcelona, http://www.telemedicineclinic.com/. Two Swedish hospitals (in Borås and Sollefteå) established a connection to TMC in March 2003 [DGINFSO06], and the TMC web site claims that the company is the only teleradiology service provider with an established connection to NHSnet in the UK and Sjunet in Sweden. There is also a report on two hospitals in Mid-

Norway (in Namsos and Levanger), that established a connection to TMC in 2004 [Bergstrøm04], but the arrangement was later terminated for non-technical reasons [Skjetne07]. According to a presentation by TMC CTO Johnny Eriksson at Röntgenveckan 2006, their client list also includes hospitals in England, Stord Hospital in Norway and several more in Sweden [Eriksson07].

An example of cross-border exchange of medical information (and indeed of human organs) in another medical genre is the Eurotransplant International Foundation (http://www.eurotransplant.nl). The organization is responsible for the mediation and allocation of organ donation procedures in Austria, Belgium, Germany, Luxemburg, the Netherlands, Slovenia and Croatia. In addition to telephone and fax, remote system access to the so-called donor procedure applications is used [Slot07].

Health-related tele-education projects across national borders are numerous. The Pacific Open Learning Health Net (http://polhn.org) may serve as an example of a rather broad curriculum through distance education in the Pacific Islands. Another example, with a considerably more specialized focus, is a series of second opinion consultations in telepathology, conducted between the universities in Teheran and Basel [Abdirad06].

3 Network performance testing methodology

3.1 Introduction

The task of network performance measurements is an activity undertaken by a wide range of actors, and with differing perspectives. Researchers on network protocols measure network performance, and use the findings to characterize and compare their implementations. Internet researchers measure performance to investigate the long-term development of international connectivity [Cottrell02]. Network architects measure performance to verify the characteristics of newly acquired network links or accesses before they are put into production use. Network managers measure performance in "live" networks to help pinpoint error situations or to assess available resources in order to plan upgrades. And network users measure network performance between end systems to find out whether the network connection is able to deliver the performance that their application requires [Uninett06]. A related activity of system reliability measurements is frequently performed concurrently, but will not be elaborated in the current context [XIWT98].

3.1.1 Metrics

Although the focus, motives and abilities of the personnel performing the measurements may differ, the performance measurement activity is in general aimed at establishing the magnitude, measured over some time period, of one or more of the network metrics throughput, loss rate, delay and delay variations. How these metrics influence real-time applications is discussed in section 4.2.3.

Without any specific application in mind, it is realistic to test for two common traffic patterns, which are shared by a number of applications. These traffic patterns also represent some of the most demanding loads on networks in terms of resource allocation and fair use. TCP [RFC793], as a connection oriented, sliding window protocol, employs flow control mechanisms to utilize the maximum bandwidth available to have an error-free copy of an amount of data transported from a sender to a receiver in the shortest possible time frame. The amount of digital data delivered per time unit is the *throughput* of the network path[4] between the two nodes, and is measured in bits per second (bps).

Another important characteristic is the *loss rate*, which is the ratio of packets sent but not received, to the total number of packets sent, for a time interval, over a network path. Data packets in transmission may be discarded for several reasons. Network congestion occurs when a transmission buffer in a router is filled to capacity, and even more data packets are being routed towards the resource (link) that the buffer is queueing for, causing the excess packets to be discarded, Bit errors in packets may be detected by checksum verification

[4] A network path is by definition unidirectional [RFC2330].

procedures in the transmission equipment or by the end system, causing packet drop. Even IP time-to-live header field decremented to zero may cause packet loss, e.g. due to IP routing loops, in which case the loss rate is most likely 100% until the situation is corrected.

It is however important to realize that a certain loss rate is normal in a IP data network: the flow control mechanism in TCP is based on detection of packet loss, and retransmission of the dropped packet(s) at a slower pace. The loss rate becomes significant in its own right when considering protocols without inherent retransmission mechanisms, e.g. UDP [RFC768]. UDP is increasingly used in conjunction with RTP [RFC1889] to transmit real time audio and video, where any loss rate significantly above 0% is experienced as a reduction in service quality. A more stringent account of a packet loss metric in IP networks is presented in [RFC2680].

Network *delay* in general refers to the one-way delay along a network path, and denotes the time interval from the first bit of the packet hits the wire, until the last bit of the packet is received. The term round-trip delay (or round-trip time, RTT) is used for the sum of the delays in both directions. The RTT is not necessarily a sum of two equal parts, as asymmetric routing will cause different paths to be traversed for data packets in each direction. Also, even when the two paths are symmetric, they may have different performance characteristics due to asymmetric queueing. A more stringent account of a round-trip delay metric in IP networks is presented in [RFC2681].

The *delay variation* is usually interpreted to be the maximum difference in delay between any two packets in a stream, in a specific interval. One important use of it is to determine the minimum size of the play-out buffer in audio and video applications. The delay variation metric is sometimes called *jitter*, although this term has a slightly different meaning in telecom than in datacom. A more stringent account of a delay variation metric in IP networks is presented in [RFC3393].

The ITU-T standard G.114 recommends 150 ms as the maximum one-way delay for "good" interactivity [ITU00]. An extension of these recommendations to include (one-way) jitter and packet loss is proposed in [Calyam05]. The ITU has also recently completed a more comprehensive work, with definitions of IP packet transfer performance parameters [ITU06a] and specification of six different QoS classes based on various IP applications [ITU06b].

Table 1 Network metric categorization [Calyam05]

	Good	Acceptable	Poor
Delay	0ms-150ms	150ms-300ms	>300ms
Jitter	0ms-20ms	20ms-50ms	>50ms
Loss	0%-0.5%	0.5%-1.5%	>1.5%

Streams of related packets may be routed across different paths, due to topological redundancy in the network. As this paths may display different delay characteristics, the ordering of the packets may be different when received compared to when they were sent. This *out-of-order delivery* rate is occasionally included as a separate QoS parameter, although the consequences and correction strategy for real-time applications are the same as for delay variation.

3.2 Measurement approach

Although the object network for the test has no official status as a transport channel for medical or other utility applications, some of its components obviously play an important role in official service delivery in the national health networks. With performance measurements there is always the risk that the injected test traffic will influence production services by offering increased competition for network resources. To minimize the adverse effect of such a potential "denial of service" situation, one should try to minimize the duration of each measurement run. For the same reason, test runs that are required to have some duration to achieve credible mean values, should be performed at hours with minimum competing demands for network bandwidth.

The measurement surroundings have posed some constraints on choice of equipment, tools and number of test sites. As it was not realistic to acquire dedicated test equipment for the experiments, the measurement activity had to be performed with what was already in use in the national health network organizations, and could be made available for the thesis work. Similar considerations regarding the range of test sites, and the observation that test execution at each site would require the participation and cooperation of a volunteering network operator, made it unrealistic to perform the measurements across a mesh of sites.

Based on availability and prior experience, the software product IxChariot [Ixia07] was selected to perform network measurements between Norsk Helsenett's branch office in Tromsø, and the Sjunet connected organization Landstinget Västmanland (LTV.se). The tools are being used in Norsk Helsenett to verify the bandwidth of IP-VPN accesses as they are delivered in the WANDA project, and have been used for similar purposes in Sjunet.

3.2.1 IxChariot operation mode

The network test application IxChariot consists of a management console executing on a Microsoft Windows workstation, and two or more test probes which may execute under Windows or a number Unix variants, or even a selection of dedicated probe hardware. Network testing scenarios are described in proprietary language scripts. IxChariot comes with a wide range of scripts, most of which are designed to emulate specific application behaviour like database transactions, DNS lookups etc. But the script portfolio also includes a selection of throughput measurement alternatives.

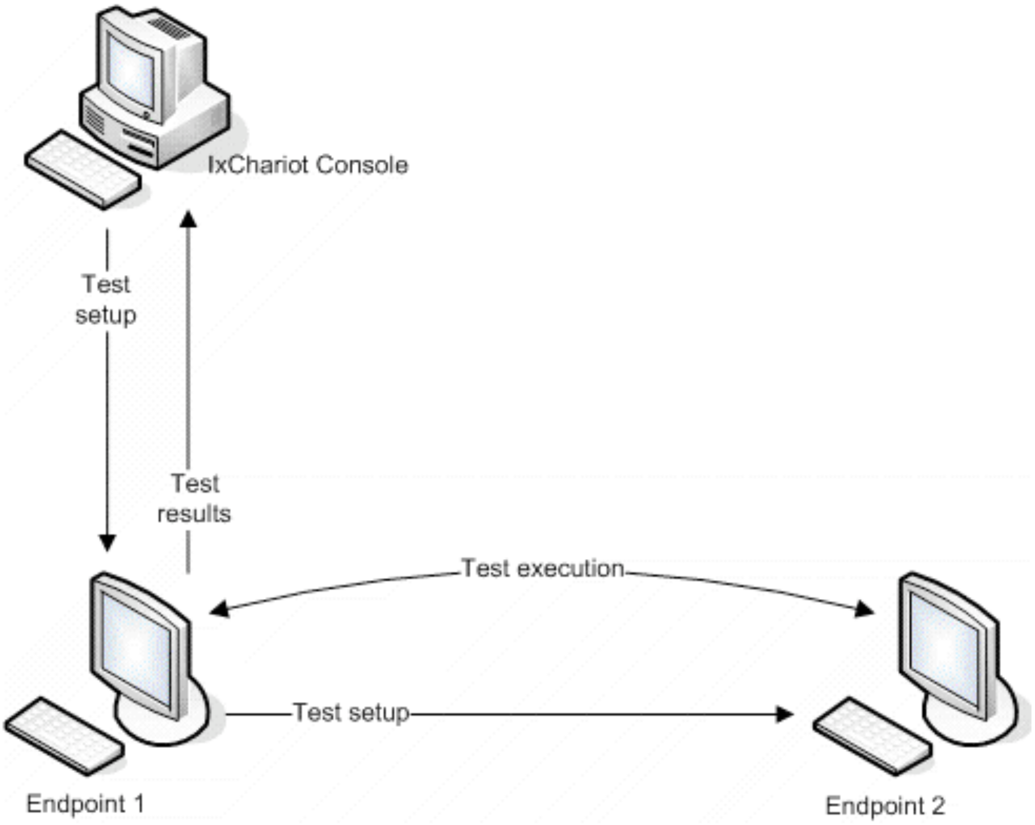

Figure 7 IxChariot test process [Ixia06]

A control script for the intended test scenario is created on the workstation and downloaded to endpoint 1 (the master probe), which in turn transfers control instructions to endpoint 2 (the slave probe). The test is executed between the two probes without intervention from the console, and the results are returned from the master probe to the console upon termination of the test execution.

3.2.2 Firewall considerations

Although the IxChariot console and one of the probes are co-located in Tromsø, the network path between the two probes traverses infrastructure in all three national health networks, each of which with firewalls in place to manage network security. Lab simulations were performed to investigate IxChariot's operational capabilities in a firewalled network environment.

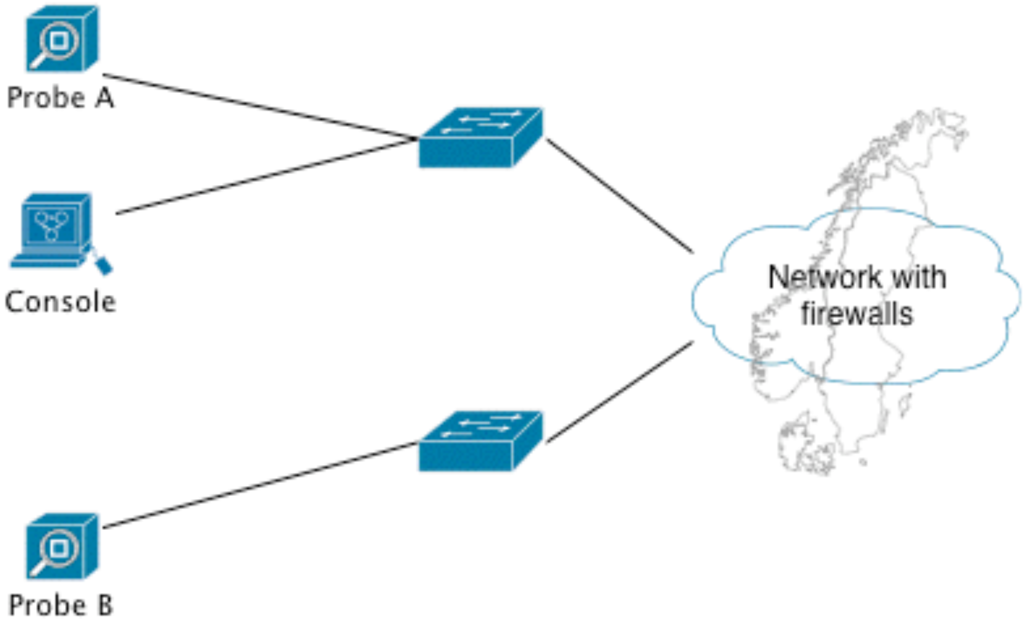

Figure 8 Firewalled test network path

The successful execution of a test scenario depends on the installation of access privileges for the relevant traffic to pass through all firewalls en route. This in turn can only be done for communication sessions with statically assigned TCP and UDP port numbers, and in each of the three traffic categories: test setup, test execution and test results reporting.

The test setup traffic uses port 10115/TCP, and 10115/UDP for jitter measurements when the RTP protocol is employed. While applications that use dynamically assigned port numbers may be a challenge to allow through firewalls, static port numbers are readily configured. The test execution traffic uses a script configurable destination port, with the same consequence. Also, when the local probe (probe A) played the role of endpoint 1 the results reporting comprised of no firewall traversals, and hence was without problems. But the test results reporting from the remote probe (probe B) turned out to be more of a problem.

There turned out to be three options. The Ixia script editor offers the option to switch the roles of the sender and the receiver. This would make the probe A initiate transfer from probe B, and have the results reporting occur locally even when the traffic load direction is B-to-A. However, this option is not available for UDP (called *streaming scripts* in the Ixia documentation). Another alternative was to instruct IxChariot to have the remote probe re-use the test setup connection for results reporting. But selecting this option in the IxChariot menu system caused the application to fail after reporting a single timing record to the console. The third option, which turned out to be successful, was to configure IxChariot to use a static port number when reporting results back to the console, and have this traffic category become “firewall friendly” as well.

3.2.3 Baselining

Prior to test execution, a lab was set up to simulate the measurement scenario. A two-interfaced server with FreeBSD 6.2 was installed, and a laptop with Linux Ubuntu 2.6.17 and IxChariot end point software connected via a switch to each of the interfaces. The server had the software modules `ipfw` [FreeBSD06] and `dummynet` [FreeBSD02] configured, which enabled it to simulate arbitrary network delay, bandwidth limitations and even packet loss for traffic passed between the two interfaces.

In addition to verifying the sensibility of the intended tests, the lab environment was useful for determining the IxChariot firewall capabilities and their consequences to network traffic, as described in the previous section. Ideally the same two probes would have been used in the real measurements as well, but only one of them could be made available for the required number of days.

3.3 Test critique

Unfortunately, it turned out to be unfeasible to conduct the performance measurements as planned within the time constraints of this thesis work. The execution of the tests presupposed configuration of equipment in the health networks of all three Scandinavian countries. One month before the thesis deadline it turned out that the resources could not be found to complete all of this configuration work in time, due to a high workload of other, more urgent projects.

As a practical alternative, it was decided to perform measurements over the Internet without use of VPNs, by configuring reciprocal NATed firewall openings at Internet access points in Sjunet and Norsk Helsenett. The purpose of such a test is two-fold. First, it can serve as a "test run" of the experiment, providing experience and possibly improved methodology when time and resources can be found to execute the measurements via the Sundheds-DIX. Second, the results can provide a baseline for result comparison with the "real" experiment.

3.3.1 IxChariot measurements to LTV, Sweden

As the main network path of this test was over the open Internet, the resulting bandwidth figures have limited relevance, apart from an indication of the Internet access capacity of the end-point networks. No mechanisms were employed to secure the confidentiality of the (test) data.

The measurements are in lack of a UDP component, as noted in section 5.2. Both this problem, and the failure to activate window scaling in the TCP measurements could have been avoided if the equipment used for lab baselining had been available longer, and could have been installed at the test site at LTV. However, the latter problem had only modest impact on the TCP measurements, as the option to utilize several parallel TCP sessions was available. The aggregated throughput of these sessions adequately answers the question of throughput along the path tested.

Regarding the lack of UDP measurements, an effort to investigate the situation in order to diagnose and possibly correct the problem was pre-empted because it proved necessary to re-allocate the other "lab probe" to production use as well. The unanswered UDP related questions, such as packet loss and delay variation could alternatively have been answered by employing other tools. But in this case, the time was not available for introduction of yet another measurement regime.

3.3.2 Iperf measurements to S-DIX, Denmark

Although firewall openings in the Sundheds-DIX could not be configured in time, it turned out that there were openings in place for measuring performance with an alternative tool, Iperf [NLANR05], as part of more general troubleshooting and support procedures [MedCom07e]. This would provide some insight into the performance of the established Scandinavian health network. However, this measurement regime has some deficiencies.

Measurements are one-way only. Firewall openings are in place to allow network traffic to a centrally located Iperf server, but corresponding firewall openings are not in place to switch roles and generate traffic in the opposite direction. This means that one can only assume without means of verification that network performance for connected entities is approximately symmetrical, unless it is known that asymmetrical technologies like ADSL is used en route. Besides, it may also give a positively biased impression of the available bandwidth, as e.g. some e-learning scenarios use a one-to-many traffic distribution model with of a centrally located MCU. In such cases, the dominant traffic load would be out from the Sundheds-DIX, i.e. in the opposite direction of what is offered for Iperf measurements. And finally, the implemented Iperf scenario doesn't enable measurements between two connected health institutions, as that would presuppose traffic initiated towards a connected entity, i.e. towards the periphery in the hub-and-spoke network topology.

Reduced resolution. While the Chariot tool has a mechanism for periodic reporting of partial results, and uses this to produce a graph of the development of a test, Iperf supports a similar mechanism for TCP only. The intermediate results reported for UDP measurements conducted with the "`--interval N`" switch are the traffic load level delivered to the network by the sending application, and not the amount received from the network by the receiving application.

Possibly **reduced throughput in individual TCP sessions**. It was observed that Iperf, when instructed to increase the window size of a TCP session by use of the "`--window N`" switch, issued a warning that the window size had been reduced to 256 kByte, while there are indications that the window size actually used does not exceed 64 kByte. Further investigations are needed to uncover the cause of this (e.g. it may be an artefact of one of the operating systems involved), but the fact remains that it was necessary to operate three TCP-sessions in parallel to achieve maximum TCP throughput to Sundheds-DIX with Iperf.

Fragile tool/platform. On the Ubuntu (Linux) platform that the measurements were performed from, individual UDP measurement often failed with one of the error messages

- `WARNING: did not receive ack of last datagram after 10 tries.`
- `read failed: Message too long`
- `write failed: Message too long`

Even worse, an aborted measurement would also imply that the following measurement had to be dismissed, as the server would report combined results for the two with no means of separating the wheat from the chaff. Overall, more than two thirds of the individual measurements were useless. However, it is appropriate to note that this situation may have other causes than the Iperf software itself (cf. [Debian06]), and is not in accordance with previous experience with the tool on FreeBSD.

Even more serious was the dramatically high packet loss figures reported by Iperf. A 384 kbps UDP stream was reported to display a packet loss rate slightly below 10%. It does not require much experience from data network operations to realize that this is not consistent with the performance figures that were achieved with 64kByte TCP window size. The conclusion is that the UDP Iperf results are unreliable, and should not be trusted.

4 Infrastructure requirement and design issues of a Scandinavian health network

A Scandinavian health network should fulfil the requirements to national health networks from section 2.5, but be implemented with the geographical perspective of Scandinavia. For practical reasons it would be advisable to reuse as much as possible from the established national networks, but the only initial restriction is that it should be functional with respect to the purpose: to serve the health sector in the Scandinavian countries.

4.1 Introduction

The operational requirements posed to a health network are a function of the services that the network is intended to carry. Although an attempt could be made to list all medical, administrative and other applications that a health network should provide transport for, such a list could not possibly be complete. The sheer number and variety of networked applications in any hospital of some size may be counted in decades. Connecting hospitals into regional and national health networks makes it relevant to provide many of these services across these networks. When the health networks even cross national borders, the size of the application mix is bound for rapid increase. Adding the time factor to this equation, and the proposition that the application portfolio is likely to change with time, is becomes clear that creating and maintaining such a list would deserve to be characterized as Sisyphus work.

4.2 Network service categories

Based on the observation above, it may be more sensible to turn the question upside-down, and categorize networked applications based on the requirements they put on the network.

4.2.1 Message oriented services

Much of the early communication activities in the health sector focused on exchange of medical data in structured email messages, later known as EDI. A Norwegian example is a trial service that emerged in May 1990 from the newly established telemedicine environment in Tromsø, consisting of transfer of laboratory results from Clinical Chemical Laboratory at RST to the municipal health service (GP) in Bardu. The service employed modems and PADs to access Telemax.400 [Gamst07], an X.400 email service offering from Televerket that was to become commercial a year later [Hausken91]. Similar activities took place during the 1990s in many countries.

Although practical benefit of EDI messages was proven early in the telemedicine history, it would be wrong to suggest that these initial services are old-fashioned, and that the technological development is in the process of making EDI based services passé. On the contrary: EDI based message exchange is a primary "work horse" in the provisioning of telemedicine services, as is convincingly illustrated by MedCom's statistics from the Danish health network [MedCom06a].

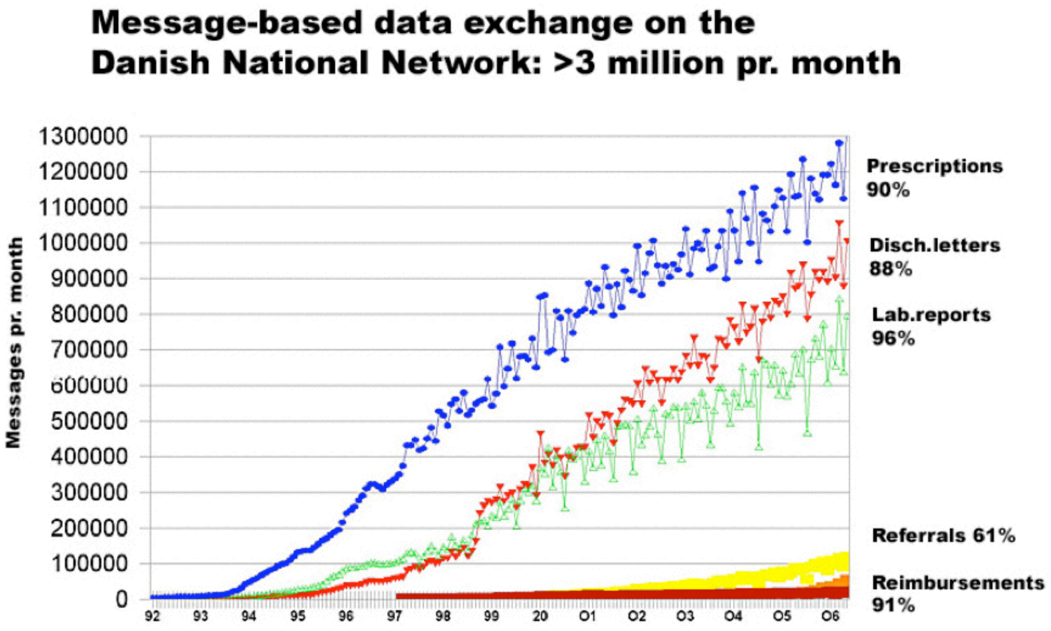

Figure 9 Monthly exchange of EDI messages in the Danish health network [Bech07]

Besides the immediate payback in terms of useful services, modest requirements to the underlying technical infrastructure makes message oriented services attractive. Email UAs (clients) may communicate with an MTA (server) with very modest requirements to the communication channel, as indicated by the Tromsø example above.

4.2.2 Interactive services

Although it is hard to imagine a health network without exchange of email messages (EDI), a health network based on EDI messages alone would be very limiting in terms of service capabilities. Paralleling the growth of world wide web on the Internet, browser based access to services have become popular in health networks as well. Web browser technology is well suited not only for the traditional navigation through hyper text document space, but is frequently adapted as front-end to database accessing applications – even complex applications like PACS/RIS and EPJ.

What these applications have in common with the more traditional special-purpose client/server application is the requirement of prompt response. Users of interactive applications expect to see feedback to their input within some reasonable time frame. The duration of the acceptable waiting time is no uniform entity: it may range from sub-second response time expectation in remote terminal applications, to several seconds time lag after completing a web form to request a more comprehensive database operation. Although the term "IP-over-email" occasionally surfaces in discussions of protocol layering, and some even claim to have implemented it [Ranum99], the packet delay considerations of the concept may serve to illustrate how the email platform is bound to fail as a basis for providing services with interactive characteristics.

The shortcomings of message based infrastructure in providing the required variety of services can also be illustrated by a development step of the Danish health network. It was not until MedCom IV in 2003, and introduction of the Sundheds-DIX [MedCom04], that the Danish health network could offer interactive capabilities to health organizations in Denmark. One of the factors that motivated this step was a realization that the purely message transporting traditional health network was inadequate for interactive services [MedCom01].

Although the interactive capability of a network is not easy to characterize objectively, it could be made subject of practical constraints in order to make concrete requirements to the underlying technical infrastructure. An example of such a constraint could be that the network delay between any two hosts connected to the network should not exceed 100 ms, measured on an idle network. It is important to realize, however, that in line with the characterization of traditional IP technology as "best effort service", such a requirement may become void as the collective load on a network changes.

4.2.3 Real-time services

For some categories of applications, the best effort delivery of classical IP is simply not good enough. Audio is a frequently used example where "one data packet arriving too late and we hear it" [Chafe00]. In general, real-time delivery of audio and video is based on playback with the same speed in the receiving end as the sample rate at the sending end. This makes A/V services sensitive to variations in network delay, which manifest themselves in audio disturbances and picture artefacts.

The applications' traditional "self defence" towards such delay variations is buffering in the receiving end, to produce a more uniform time spacing in the data stream before it is decoded and played back to the user. The downside to this strategy is that it introduces additional delay, on top of several other delay categories: sampling/coding/packaging delay at the sending end, network transmission delay, and un-packaging/decoding/playback at the receiving end.

For applications that transmit audio/video in one direction only, providing a service similar to the broadcasted transmissions of radio and TV stations (although not necessarily with the one-to-many feature of ether media), increased delays are usually no drawback to the listener or viewer. But for conferencing applications the buffer increasing strategy may lead to unacceptable dialogue conditions.

Delay variations that exceed the correcting capabilities of the receiving buffer are as disturbing for audio/video application users as packet loss. A packet delivered so late by the network that rendering its contents would interfere with the processing of the next data packet must be dismissed by the application. So in order to provide audio/video conferencing services of predictable quality in a health network, all the QoS parameters bandwidth, packet loss, delay and delay variation must be guaranteed by the network to stay within agreed limits. An account of how a video conferencing application may suffer from lack of QoS is presented in [Høykom03].

Although telephony and videoconferencing are the most wide spread real-time network services, more spectacular uses have been demonstrated. On September 7. 2001, the first fully remote operation ("The Lindbergh Operation") was performed by a team led by Professor Jacques Marescaux in New York, performing laparoscopic cholecystectomy on a French patient in Strasbourg [Chall03]. And only a year and a half later the first telerobotic remote surgical service was established at St. Joseph's Hospital, Hamilton, Ontario, Canada [Anvari05]. Even though TCP is obviously best suited for remote control of the instruments in such applications, UDP would be used for the audio/video feedback [LeParc02].

In addition to its application in remote operations, robotic surgery displays other features that have motivated its introduction even in local surgery at some Swedish hospitals. Four units are currently operational and more are being considered. One of the proponents, Professor Peter Wiklund at The Karolinska University Hospital, lists the benefits: "…that the operator views a very enlarged and three-dimensional picture of the operation wound; that there is a downscaling between the operator's hand movements and the instruments inside the patient which cause shakings to be minimized. Additionally one has, in contrast to ordinary laparoscopy, a «wrist» or «EndoWrist» that offers the surgeon improved maneuverability." [Ramel06] An increase in acquisitions of robotic surgery equipment for local use may contribute to lowering the barriers against using the same type of equipment for remote operations as well.

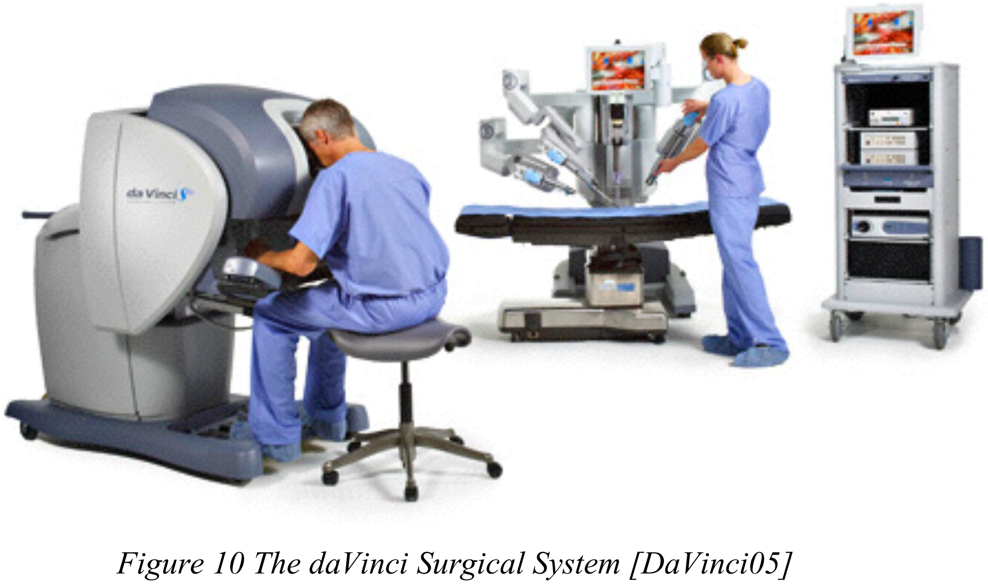

Figure 10 The daVinci Surgical System [DaVinci05]

4.2.4 Basic infrastructure requirements

Requirement Infra#1:
SHN should be able to transport real-time services between connected organizations.
Fit criterion :
Delay, delay variation and loss in SHN must not exceed the level marked "acceptable" in Table 1 (p.36) for relevant traffic.

This requirement is parallel to requirements to a future Sjunet in [Arvidsson06], and to Wanda/Norsk Helsenett in [NHNWAN05]. A consequence of the requirement is that VPNs over Internet is not a viable transport technology platform. VPNs over Internet is also dismissed as an alternative in [Haug97] and [Arvidsson06].

Requirement Infra#2:
SHN should provide non-bureaucratic technical arrangements for traffic exchange agreements and implementation.
Fit criterion:
SHN must not contain internal firewalls.

There are good reasons for any organization to protect its assets by employing security barriers, frequently referred to as firewalls. This also holds for health institutions with a requirement to regulate how their systems and data may be accessed from other members of the health sector. But sometimes motivation can be found to introduce additional firewalls on "natural" borders in a network. The rationale of such installations should be carefully scrutinized, as their operational consequence is an additional level of agreements and system configuration necessary to allow the network to fulfil its intention: communication between consenting parties in accordance with legislative requirements.

Section 6.1 (p.61) presents an account on firewalls' tendency to proliferate.

4.3 IP addressing and NAT

IP addresses have a dual functionality, serving both as identifiers and locators. The identifier property is used throughout a communication session to identify each of the hosts towards the other, while the locator is used to locate the place in the network topology where the destination host is attached. Both functionalities have uniqueness requirements [RFC2101].

4.3.1 IPv4 addresses as a limited resource

The claim that the Internet is running out on IPv4 addresses is not a new one. However, some forecasts are better founded than others, as e.g. Geoff Huston's projections at Potaroo [Potaroo07]. The prognosis here is that the IPv4 address pool will be exhausted some time in 2010.

This by no means implies the end of the Internet, or that the current IPv4 infrastructure will be useless from that date. It means that the option of obtaining additional public IPv4 addresses for new applications and connections will be exhausted. One of the most plausible consequences of this is an increased focus and activity in the IPv6 arena, with enterprises commencing a transition process. For many, even in the health sector, this process has already begun. E.g. with NHS, that has included a requirement to its N3 service provider "to develop a strategy to support IPv6 within the N3 network in the future" [Divaharan05].

4.3.2 Elimination of NAT

The problematic aspects of NAT are presented in section 2.1.2. To abolish NAT internally it is required to establish a coordinated IP network wide addressing regime, where each host has its unique IP address. Three alternatives exist for a NAT-free network:

1) public IPv4 addresses
2) private IPv4 addresses [RFC1918]
3) IPv6 addresses

To get a (very) rough sense of scale of the IPv4 address space that would be needed, one could make a comparison to the academic sector. The Norwegian academic network Uninett's AS 224 exports the equivalent of 17 /16 IPv4 networks [Yorine07]. Assuming similar figures for Sunet and Forskningsnettet, adjusted for the population distribution among the Scandinavian countries, these three networks have a total address space equivalent of a /10 network range. Further assuming commensurable size with the health sector, measured in number of networked hosts, and considering the current state of affairs with respect to IPv4 address space exhaustion, the conclusion is that obtaining a sufficient range of public IPv4 addresses from RIPE, for a wholesale conversion of all health institutions in Scandinavia seems futile.

The second option, to convert the entire network to IPv4 private addresses may seem slightly more plausible. This was the choice for the third generation NHS network, N3, completed in March 2007 [Divaharan05]. The number of health institutions that would have to renumber their internal networks under a similar Scandinavian scheme is hard to assess, but if we assume that the entire private address range from [RFC1918] is completely allocated in each of the three countries, with no domestic overlap, then two third of the hospitals and GPs would be required to reconfigure all their computers and network equipment.

And even if successfully completed, such an operation could not guarantee a permanently NAT-free network environment, as it might be expected that requirements for new connections to additional municipal, regional and national health networks will surface in the future. The experience from Norsk Helsenett is that this aspect of IPv4 private address use is problematic, as can be illustrated by one of the conclusions from the Ses@m project, connecting units in Tromsø municipality to Norsk Helsenett [Abelsen06]:

> "[it is] in general very unfortunate that NHN uses private IP addresses in their network".

The third option, converting to IPv6 addresses could be characterized as "the future-proof alternative". But it is certainly also the most demanding in terms of resource requirements. 100% of all hosts and communication equipment will have to be reconfigured. All server applications must be examined and adjusted – some may even be without available (vendor with) source code. Very few of the technical personnel have any experience with IPv6 and will need to update their competence. The list of expensive and time-consuming tasks is long.

On the other hand, there are very few experts in Internet technology that question the proposition that a conversion to IPv6 addresses will be forced some time in the future. The question then becomes: when should such a conversion be commenced? The current status of IP addressing in the health networks makes it sensible to consider answering this question with a fixed date in the near future.

4.3.3 IP / NAT requirements

Requirement IP#1:
SHN should be enable transport of all kinds of services based on Internet technology end-to-end, modulo security motivated access restrictions.
Fit criterion:
NAT must not be employed in SHN.

Requirement IP#2:
IP addressing in SHN should be future-proof.
Fit criterion:
SHN must transport IPv6 natively.

Even for IPv6 addresses several different strategies/address ranges to choose from, each with different characteristics. Although NAT functionality is not included in IPv6, the architecture does have the concept of "private" addresses, in the sense that these addresses by design are not routed on the Internet. The Unique Local Addresses (ULA, [RFC4193]) are /48 free-for-use, address blocks and come with a randomizing algorithm designed to minimize the risk that two organizations should choose the same prefix out of the 2^40 that are available.

Considering "the birthday paradox", and assuming a truly random distribution, the probability that two prefixes are equal (causing address collision) in a network connecting N ULA subnets is given by the formula:

$$1-\frac{2^{40}!}{2^{(40*N)}*(2^{40}-N)!}$$

Table 2 Birthday paradox values for selections of N out of 2^40

N	p(N)
10	4.09 * 10^-11
10^2	4.50 * 10^-9
10^3	4.54 * 10^-7
10^4	4.55 * 10^-5
10^5	4.54 * 10^-3
10^6	3.65 * 10^-1

Table 2 shows that a health network connecting 100,000 hospitals that use ULA addresses will have a 0.5% probability of an address collision, while for 1,000,000 units the collision probability is 36.5%. If even these figures are regarded as too high, a recent initiative from APNIC could be considered. The suggestion is to re-vitalize an earlier proposal for a similar scheme with central coordination of assigned prefixes (ULA-C, [Hinden05]). If this proposal is carried through, global uniqueness of “private” IPv6 addresses will be guaranteed.

4.4 DNS

Among the differences noted in section 2.4.5 concerning the national health networks in the Scandinavian countries is the choice of DNS strategy. Both the Danish and the Norwegian DNS implementations have problems.

4.4.1 Split horizon DNS

The Danish SDN’s DNS implementation requires that resources intended to be available network-wide are registered under the artificial top level domain .medcom. This is a health network internal name space that lists the resources’ SDN-public, but non-Internet visible, IP addresses (post-NAT). Consider a hospital that has a requirement to make one of its internal medical systems available to selected cooperating institutions. The system’s private IP address is mapped to an SDN-public address through a static NAT entry in a network device. This address is subsequently associated with a name that is registered with an A RR as *something*.medcom in the SDN DNS. The paradox is that this resource is not going to be available under that name from the hospital’s internal clients, because the IP address that the name resolves to is different from the IP address that the resource is available as internally. In other words, the resource owner cannot access the resource in the same manner as the community that has been granted access by the owner.

Sjunet suffers from the same internal split horizon problem, although here some of the effect is leveraged by recommendation to make the organization’s sjunet.org domain be a copy of the organization’s official domain. Even so, the recent IBM evaluation of Sjunet presents a user requirement to have “a united DNS structure, where any system always has the same logical name in DNS, independent of whether one tries to access the system from the same network or from another one" [Arvidsson06]. A similar recommendation is given in Uninett's report on establishing a national health network in Norway [Haug97], and in the technical recommendations for the Baltic eHealth project [Nohr06].

It should be noted, however, that recommendations to the contrary can also be found. The US NIST's "Guidelines on Firewalls and Firewall Policy" is perhaps the most prominent, where security considerations is the prime motivation [Wack02]:

> An organization should maintain separate internal and external domain name servers. This practice, known as split DNS, ensures that private internal systems are never identified to persons external to the organization.

4.4.2 Common DNS root

Norsk Helsenett's DNS system suffers from an even more serious problem, in that the domains do not share a common root. Connected institutions are allowed to use their Internet domain name inside the health network as well, resulting in a complex configuration of the internal name space. Although most organizations have opted for sub-domains of nhn.no, there are also many who use second-level domains – mostly under .no, but .com, .net, .org and .nu are also represented.

The common root feature of the health network's internal DNS becomes important when connecting to other networks. In order to form a common internal DNS system and make services available across the interconnected networks, they must import each other's DNS name spaces, through query forwarding or zone transfers. If this space is fragmented, the integration task involves more configuration work. If in addition it is dynamic, as is the case with Norsk Helsenett, then the task becomes a consistency challenge. The common root feature is also among the recommended features for a national health networks in [Haug97].

4.4.3 DNS requirements

Requirement DNS#1:
The DNS service should be scalable.
Fit criterion alt. 1:
There should be a minimum number of DNS roots represented.
Fit criterion alt. 2:
The Internet DNS must be used, also for health network internal resources.

Concerning the number of common DNS roots, only the uniqueness property is actually important for scalability purposes. The ability to access the entire health network DNS name space is vital to enable users' access to services across institutions. The more fragmented and dynamic this name space is, the higher the risk that only an incomplete subset of the name space will be available in some parts of the network.

The absolute minimum number of DNS roots is of course one, meaning that all health network resources share a common DNS suffix. The current Scandinavian health network, connecting the Swedish and Norwegian health networks to the Danish S-DIX, uses this approach with the .medcom domain. If a more neutral and descriptive name (like .health) could be registered with ICANN, this could possibly be turned into a viable strategy – possibly with ISO-3166 two-letter country codes as second-level sub-domains. An obvious drawback is that such a renaming of the entire health sector's network appearance would be expensive, to say the least.

For the other alternative, there are two important aspects to consider before deciding to use the Internet DNS as the only DNS system in the health network. One is the desirability of making all resource names publicly known, and the second is the choice of IP addressing strategy for the network. The security implications of a publicly known DNS name space are bound to be subject to considerable controversy. While e.g. the US NIST argue that the requirement to hide internal DNS data is strong enough to warrant a split DNS, others claim that this is argument belongs in the security-through-obscurity category, with minimal valid security relevance.

Regarding choice of IP addressing strategy, certain IP addresses preclude registration in the Internet DNS. For IPv4 private addresses, [RFC1918] mandates that DNS RRs and other information referring to internal private addresses should be contained within the enterprise. For IPv6 ULA addresses, [RFC4193] states that nodes with only local IPv6 addresses must not be installed in the global DNS. The motivation behind both of these restrictions is the lack of global uniqueness for the address categories involved. The consequence is that these address types can only be registered in an internal DNS, requiring a split-horizon DNS installation. Consequently, a split-horizon DNS can only be avoided by using IPv4 public addresses, IPv6 global unicast addresses or IPv6 ULA-C addresses.

5 Test results

5.1 Introduction

As the test were planned to be performed on the network described in section 2.4.6, the major part of the network paths involved crosses the Internet. This adds a level of unpredictability to the results, as the load conditions may change at random. Also, it is difficult to assess the causes of the throughput limits uncovered – the Internet access capacity of the networks involved, bandwidth related encapsulation limitations on the tunnel endpoints, or other restrictions or limitations along the network paths.

5.2 IxChariot / StensPC.ltv.se

A network trace between one of Norsk Helsenett's Internet access points in Tromsø and Landstinget Västmanland includes 10 hops before the traffic gets neglected by the firewall in hop 11. Uninett (AS224) is traversed in hops 2-4, NORDUnet (AS2603) in hops 5-6, before traffic is handed over to TDC/Song at the Stockholm-A peering mesh in hop 7, and transported inside their autonomous routing system (AS3246) the rest of the way.

```
: anders@bossa; traceroute 193.180.9.28
traceroute to 193.180.9.28 (193.180.9.28), 64 hops max, 44 byte packets
 1  ext-gw.rito.no (193.157.64.1)  0.375 ms  0.281 ms  0.241 ms
 2  tromso-gw.uninett.no (158.39.52.1)  1.112 ms  1.633 ms  0.637 ms
 3  trd-gw.uninett.no (128.39.47.97)  14.256 ms  14.508 ms  14.131 ms
 4  oslo-gw1.uninett.no (128.39.46.1)  24.015 ms  22.149 ms  22.389 ms
 5  no-gw.nordu.net (193.10.68.101)  22.162 ms  22.017 ms  23.330 ms
 6  se-tug.nordu.net (193.10.68.29)  29.905 ms  29.756 ms  29.835 ms
 7  netnod-ix-ge-a-sth-4470.se.sn.net (195.245.240.41)  30.612 ms  34.823 ms  30.870 ms
 8  rif10-rs1-t4-sto.se.sn.net (81.216.0.138)  31.156 ms  30.392 ms  30.810 ms
 9  rif2-cr1-vf-oby.se.sn.net (213.187.195.94)  31.432 ms  30.944 ms  30.958 ms
10  rif2-cr1-vf-vst.se.sn.net (213.187.195.98)  32.625 ms  32.104 ms  32.745 ms
11  * * *
(etc)
```

Figure 11 Network trace between NO-SE

Interestingly, a network trace in the opposite direction uncovers a routing asymmetry. Hops 7 in Figure 11 and 012 in Figure 12 are both on "Stockholm-A GigE", and marks one of the departure points. The other departure point is hop 8, which corresponds to hop 008. The asymmetry consists of the extra hops 009, 010 and 011 that the traffic from Sweden to Norway traverses.

```
Node  Time (ms)   Address           Name
(initial LTV internal hops LTV removed)
 007   2          213.187.195.97    rif3-cr1-vf-oby.se.sn.net
 008   34         213.187.195.93    rif47-rs1-t4-sto.se.sn.net
 009   2          81.216.0.137      rif5-cr3-kst-sto.se.sn.net
 010   2          213.50.65.49      ae1.kst-p1.sto.se.sn.net
 011   3          88.131.143.64     static-88.131.143.64.addr.tdcsong.se
 012   7          195.245.240.24    ne-gw-4470.nordu.net
 013   11         193.10.68.30      no-gw.nordu.net
 014   11         193.10.68.102     oslo-gw1.uninett.no
```

```
015   19         128.39.46.2        trd-gw.uninett.no
016   36         128.39.47.98       tromso-gw.uninett.no
017   33         193.157.64.1       ext-gw.rito.no
```

Figure 12 Network trace between SE-NO

The extra hops are all inside TDC/Song's network. But as these hops contribute with a maximum of five milliseconds (15%) increased round-trip time, their effect on the performance measurements may be assumed to be small.

The round-trip time measured between the two IxChariot probes was just below 35 milliseconds (Figure 13). A 33.9 ms RTT would imply that the TCP window can be transmitted up to 29.5 times each second, giving an upper bound on single-session TCP throughput of 15.1 kbps with a TCP window of 64 kByte.

```
--- 193.180.9.198 ping statistics ---
60 packets transmitted, 60 received, 0% packet loss, time 59002ms
rtt min/avg/max/mdev = 33.428/33.945/34.466/0.235 ms
```

Figure 13 Network tround-trip time between SE-NO

5.2.1 TCP throughput

An attempt to employ the mechanisms described in [RFC1323] to increase the TCP window size beyond the default 64 kByte failed, with the following error message presented on the IxChariot console:

```
CHR0125: Endpoint 2 does not support the following function(s) required to run
this test: Setting connections send and receive buffer size IPv4.
```

With IxChariot, it was possible to measure performance in both directions, even simultaneously. The achieved throughput in each of the directions was 13.3 Mbps and 13.4 Mbps, approximately 88% of the maximum estimated from the mean round-trip time.

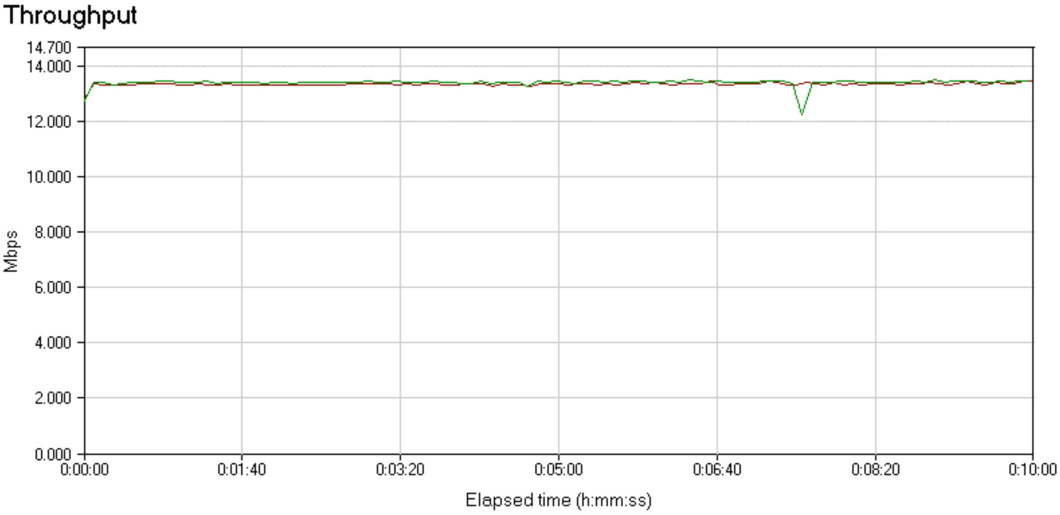

Figure 14 TCP throughput measured concurrently from and to Sjunet (LTV), 64 kB TCP window size

An assessment was made of the total throughput in both directions by running several TCP sessions in parallel, all with a window size of 64 kByte. In the direction Norway-to-Sweden, the ten-minute mean of the aggregated thoughput was 45.1 Mbps, while the Sweden-to-Norway throughput clocked in at 48.2 Mbps.

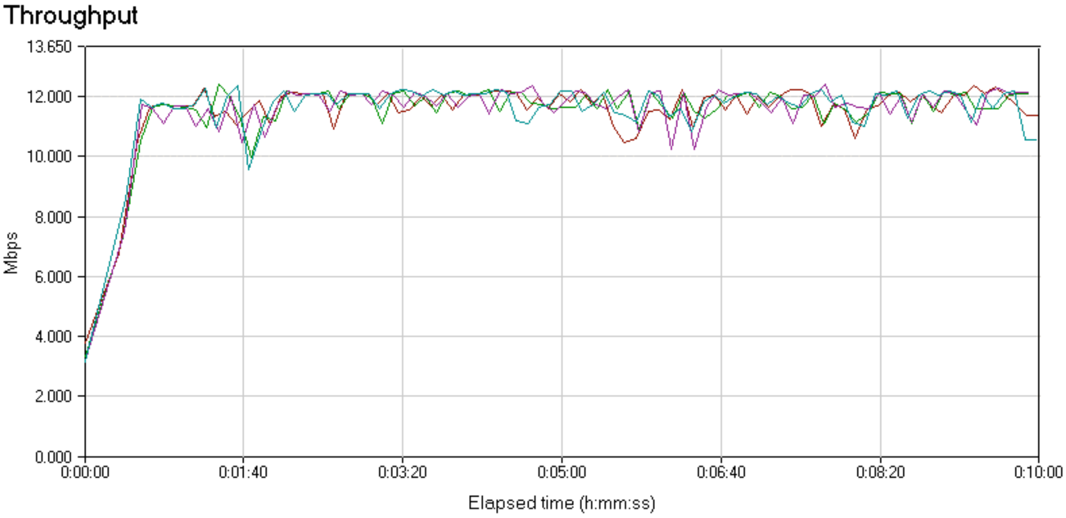

Figure 15 Four-session TCP throughput to Sjunet (LTV)

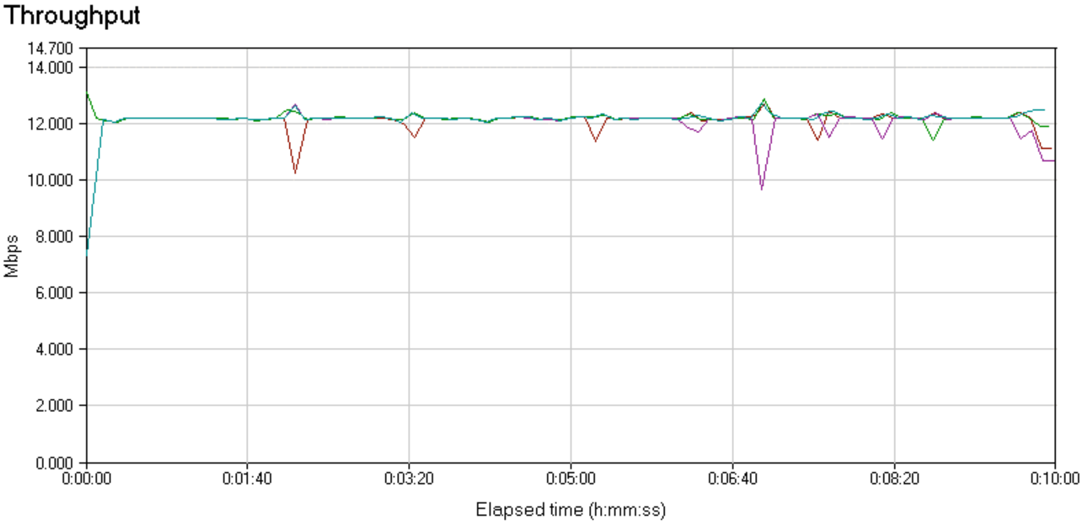

Figure 16 Four-session TCP throughput from Sjunet (LTV)

5.2.2 UDP measurements

It turned out that the planned UDP measurements, intended to investigate packet loss and delay variations, could not be performed successfully. The reason for this failure is uncertain, as it was necessary to re-allocate the local test probe to production use before further investigations and diagnosis could be executed.

The problem symptoms observed was as follows: streaming measurements could be configured and commenced as normal, but progress stopped after 10-15 seconds, and the IxChariot “Timing Records Completed” counter ceased to be incremented after just a handful of records. These measurement instances did not progress further, and had to be terminated through the “Abandon Run” functionality of IxChariot.

5.3 Iperf / falcon.uni-c.medcom

A network trace between the tunnel end points uncovered that the path includes 17 hops, and traverses Uninett (AS224) in hops 2-5, is handed over to Swipnet/Tele2 at NIX2 in Oslo (hop 6), and is transported inside their autonomous routing system (AS1257) the rest of the way.

```
: anders@bossa; traceroute gw2.sdn.uni-c.dk.
traceroute to gw2.sdn.uni-c.dk (130.228.4.254), 64 hops max, 44 byte packets
 1  ext-gw.rito.no (193.157.64.1)  0.303 ms  0.593 ms  0.252 ms
 2  tromso-gw.uninett.no (158.39.52.1)  0.328 ms  0.338 ms  0.343 ms
 3  trd-gw.uninett.no (128.39.47.97)  14.133 ms  14.195 ms  13.981 ms
 4  oslo-gw1.uninett.no (128.39.46.1)  21.786 ms  22.393 ms  21.816 ms
 5  stolav-gw.uninett.no (128.39.46.250)  21.941 ms  22.298 ms  22.377 ms
 6  193.156.120.4 (193.156.120.4)  22.165 ms  23.105 ms  22.400 ms
 7  lba-core-1.pos3-2.swip.net (130.244.192.45)  58.801 ms  199.161 ms  103.772 ms
 8  kst-core-1.pos10-0-0.swip.net (130.244.218.158)  34.777 ms  34.742 ms  34.688 ms
 9  avk-core-1.gigabiteth14-0-0.swip.net (130.244.195.161) 35.205 ms 35.236 ms 34.884 ms
10  lim-core-1.pos0-0-0.swip.net (130.244.52.162)  34.970 ms  34.713 ms  34.701 ms
11  cop1-core.pos3-0.swip.net (130.244.206.58)  35.369 ms  34.856 ms  35.328 ms
12  pos4-0.val1-core.dk.tele2.net (130.227.2.81)  34.726 ms  34.626 ms  34.805 ms
13  srp9-0.val2-core.dk.tele2.net (130.227.247.50)  35.080 ms  34.897 ms  34.917 ms
14  ge49.val-srv2b-core.dk.tele2.net (130.227.247.38)  49.236 ms  34.725 ms  34.889 ms
15  129.142.249.234 (129.142.249.234)  35.377 ms  35.165 ms  35.291 ms
16  knet-uni2-lgb.uni-c.dk (130.228.5.1)  35.189 ms  35.740 ms  35.303 ms
17  gw2.sdn.uni-c.dk (130.228.4.254)  36.617 ms *  35.783 ms
```

Figure 17 Network trace between tunnel endpoints NO-DK

For the Iperf traffic, the number of hops is only 4, with the tunnel between entries 2 and 3.

```
: anders@atp; traceroute falcon.uni-c.medcom.
traceroute to falcon.uni-c.medcom (195.80.240.112), 64 hops max, 40 byte packets
 1  g0-2-5.tromsoc1-gw.rtr.nhn.no (172.21.8.1)  3.550 ms  0.857 ms  0.799 ms
 2  f0-1-1.next-gw.rtr.nhn.no (172.21.4.157)  1.695 ms  2.835 ms  1.506 ms
 3  * * *
 4  195.80.240.112 (195.80.240.112)  39.219 ms  38.580 ms  40.349 ms
```

Figure 18 Network trace between probes NO-DK

The round-trip time measured between the Iperf client and server was just below 40 milliseconds (Figure 19). A 38.7 ms RTT would imply that the TCP window can be transmitted up to 25.9 times each second, giving an upper bound on single-session TCP throughput of 13.3 kbps with a TCP window of 64 kByte.

```
--- 195.80.240.112 ping statistics ---
60 packets transmitted, 60 received, 0% packet loss, time 59002ms
rtt min/avg/max/mdev = 37.980/38.730/39.860/0.413 ms
```

Figure 19 Network round-trip time between NO-DK

5.3.1 TCP throughput

Initial investigations (with short-running tests) of the endpoints' capabilities indicated that TCP window scaling [RFC1323] was not available on the Iperf server at MedCom. Using the Iperf parameter –w to adjust TCP window size did not improve throughput beyond a setting of 64 kByte (Figure 20), which is TCP's maximum window size without the Van Jacobson extensions described in the RFC. The actual performance measured for a single session with 64 kByte window size was 11.9 kbps, 89% of the maximum estimated from the mean round-trip time.

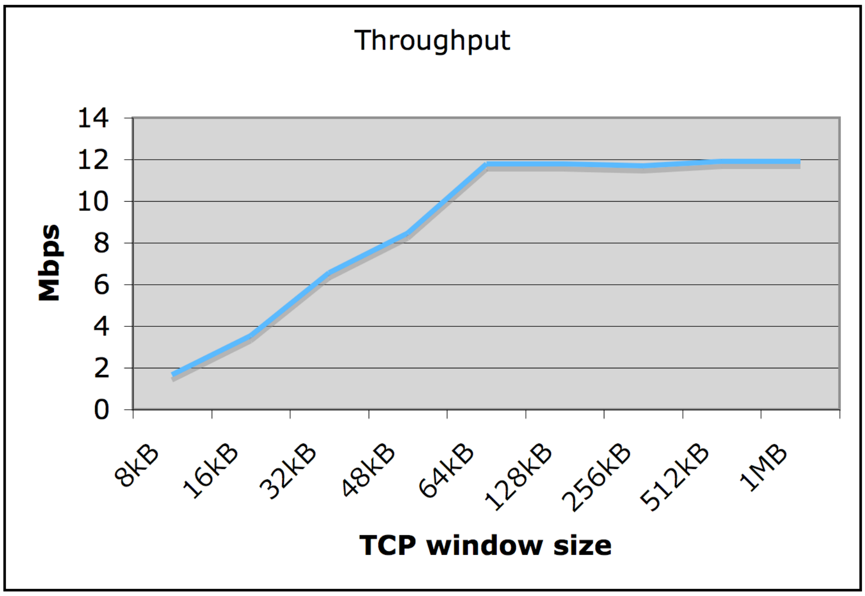

Figure 20 Single TCP session throughput with varying window size

Longer-running multi-session tests were used to assess the maximum aggregate throughput available between the end-systems. The conclusion is that there was no performance enhancement above three parallel TCP sessions, and hardly even from two to three sessions. Maximum throughput was measured to 24.2 Mbps over a ten-minute interval, with one-minute peaks as high as 24.5 Mbps (Figure 21).

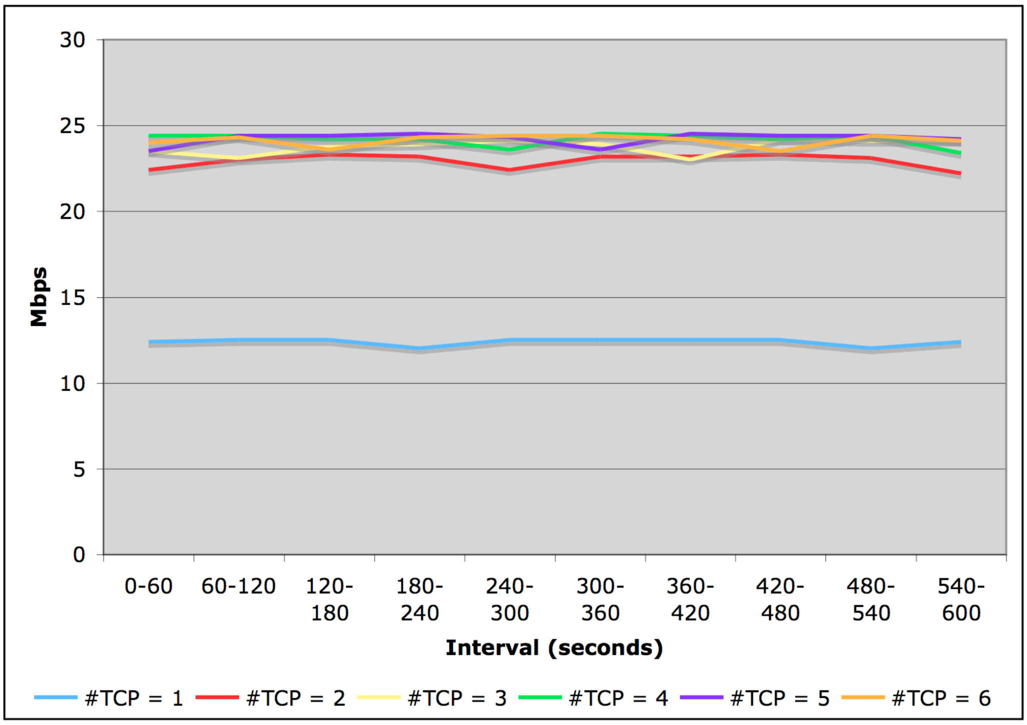

Figure 21 Aggregate throughput from multiple TCP sessions to S-DIX HUB

The three-session aggregate may serve as an illustration of variation in throughput between the individual TCP sessions (Figure 22).

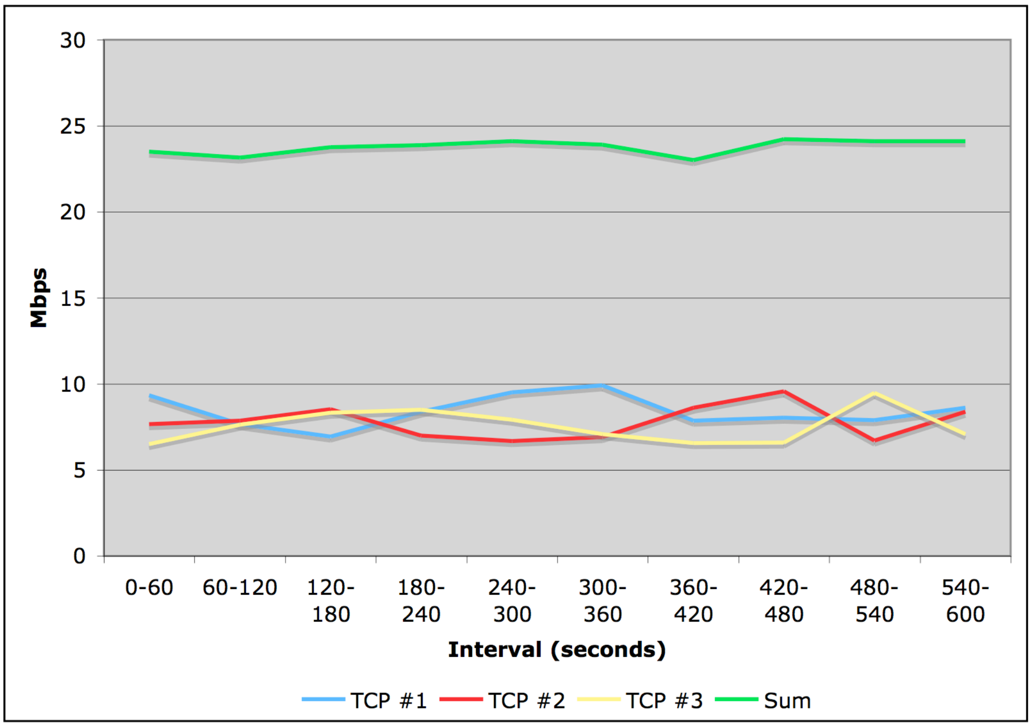

Figure 22 Three-session TCP throughput to S-DIX HUB

5.3.2 UDP packet loss and delay variation

The Iperf UDP based measurements suffered from the problem with reduced resolution mentioned in section 3.3.2. In order to produce a measurement series consisting of ten intervals it was necessary to run ten individual test series. Combined with the software/platform instability (also mentioned above), each interval in a series would be arbitrarily separated in wall time. The implication is that the graphs give a false impression of continuity in each series.

Even more serious is the obvious inconsistency with respect to the TCP bandwidth measurements. A network with a consistent loss rate as illustrated below (Figure 23) couldn't possibly deliver in excess of 12 Mbps throughput in a single TCP session. The consequence is that the measurement tool cannot be trusted – a conclusion that rubs off on the results regarding delay variations as well (Figure 24).

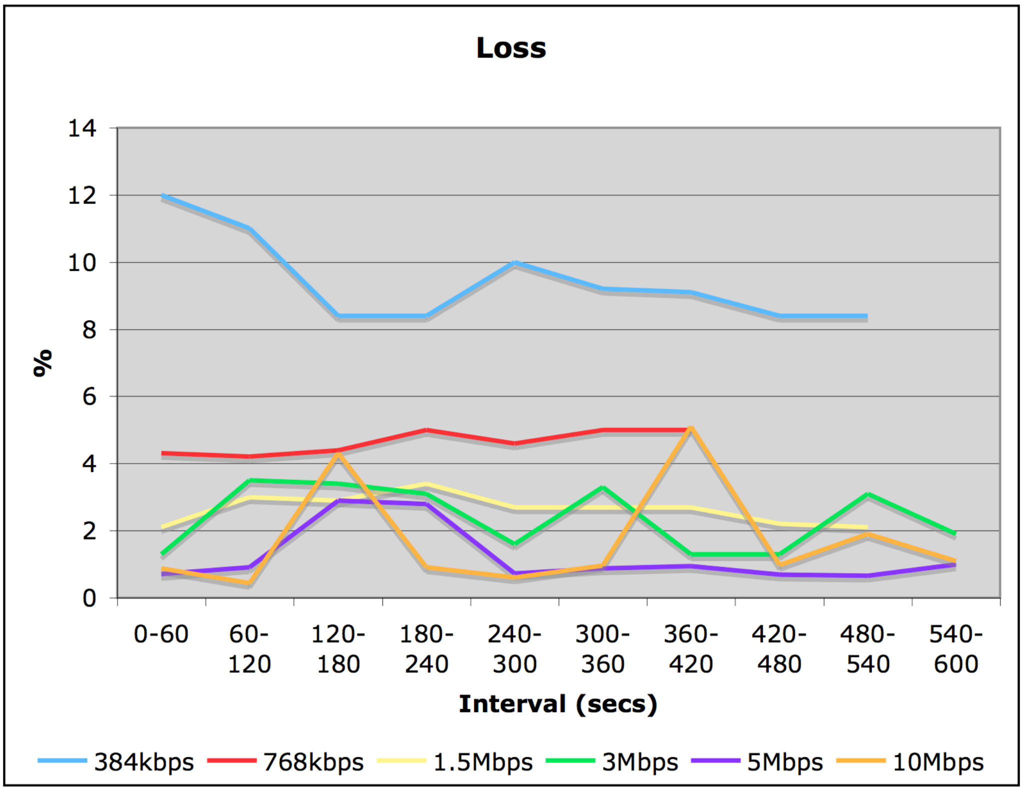

Figure 23 UDP packet loss to S-DIX HUB

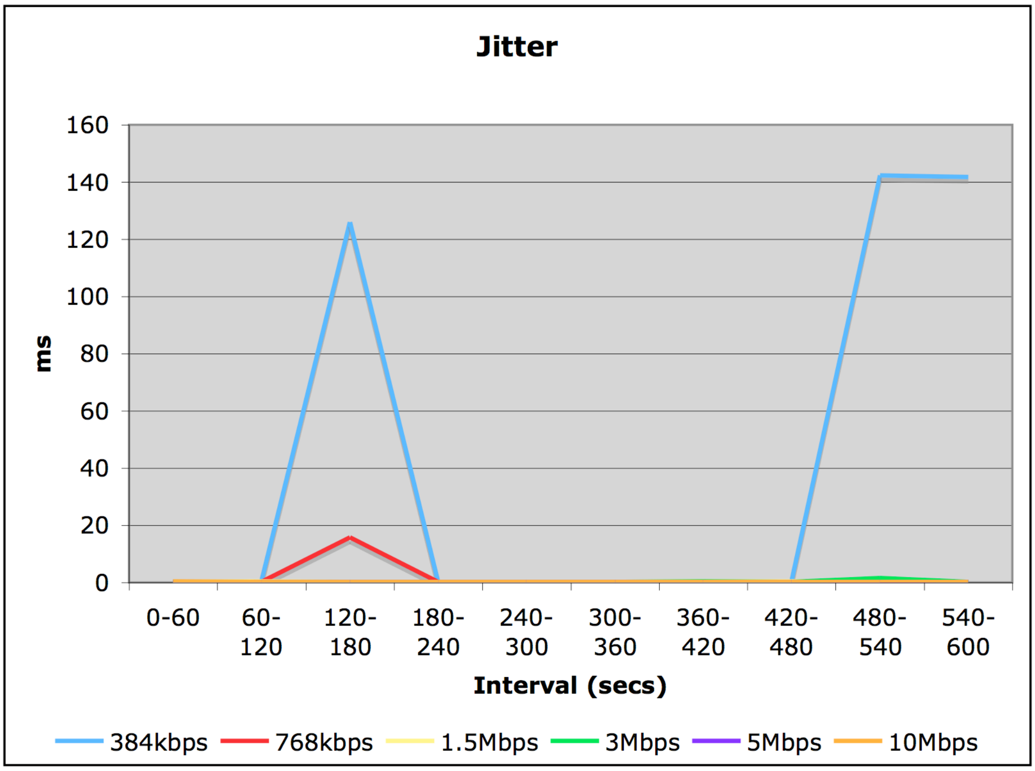

Figure 24 UDP delay variations to S-DIX HUB

6 Discussion

As noted in section 2.4.5, the national health networks in the Scandinavian countries differ on a number of issues. Some of these issues may have consequences in an effort to create a Scandinavian health network. A more detailed account of the technological issues (section 6.2) is given in chapter 4.

6.1 Structural issues

While Norsk Helsenett is a unitary network that individual health *organizations* connect to, the Swedish and Danish counterparts are networks that connect *networks* that health organizations connect to. One lesson from the Norwegian project National Health Network/Central Infrastructure, is that each connected organization tends to act as a unit and implement protective measures towards the outside world. This also holds true for organizations that are (regional) health network operators.

The goal of the project was originally to establish a high-bandwidth, high-reliability core infrastructure that all the regional health networks should connect to, in order to establish a national health network. In other words, to create a sixth network to interconnect the five that already existed. Not without similarities with the current situation in Sweden and Denmark.

The risk of introducing an excessive number of firewalls on any inter-regional network path was addressed, and firewalls were acquired in each of the five POPs in order to implement a centrally managed protective regime. This regime was administered by a group of representatives from each of the regional health networks[5]. One of the group's tasks was to implement firewalls that would protect each of the five regional networks (and their connected institutions) from the other four regional networks (and their connected institutions). It didn't work.

When the regional network organizations had their connection to the Central Infrastructure installed, they immediately acted to protect their assets by installing a regional firewall back-to-back to the POP firewall. Realizing this, the POP firewalls were soon disabled, in order to avoid yet another level of firewalls on any inter-regional network path. As a consequence, five pairs of Cisco PIX-535 firewalls were left to idle, as they became doubly redundant. The lesson to be learnt is that the number of organizational boundaries crossed by data traffic transmitted between end-systems in two connected health institutions should be given careful consideration.

The current interconnect between the national health networks in Scandinavia,

[5] The Operations & Security group, or DoS group according to its Norwegian acronym, managed by the author

via the Danish S-DIX, has some unpleasant similarities with the fragmented and regionalized infrastructure from the National Health Network period in Norway. In particular, the number of firewalls en route for data traffic between two hospitals in e.g. Norway and Sweden, and the set of organizations involved in their reconfiguration, makes the process of establishing traffic flow very complex. A symptom of this (although not odd when ownership of equipment etc is taken into account) is the fact that under the current regime, bilateral communication between institutions connected to the Swedish and Norwegian health networks cannot be configured without involving technical personnel that operate the Danish health network.

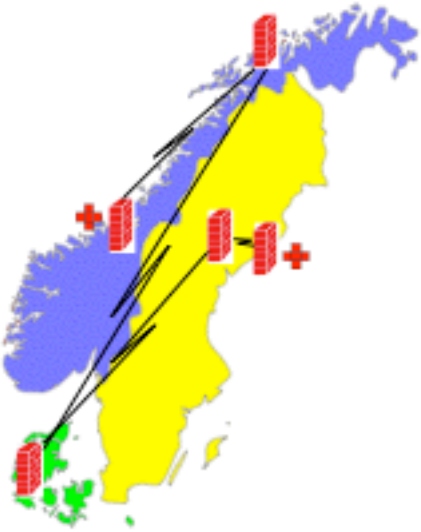

Figure 25 Path with firewalls between Norwegian and Swedish hospitals

In future editions of a Scandinavian health network, the issue of coordinated access control mechanisms should be addressed. Reducing the number of firewalls en route between any two connected health institutions to the minimum of two (one at each institution's network connection point) is an important goal for practical manageability. However, this does not in principle preclude use of centralized, common access control devices like the one implemented in the Danish S-DIX, as long as such devices are collapsed implementations of aggregated institution specific firewalls, and not an addition.

The question of admission policy in a Scandinavian health network is prudent. To what extent should categories of admissible health institution be harmonized? Should practitioners outside of the traditional medical branches be allowed to connect – chiropractors? acupuncturists? homeopaths? healers? Who should decide? And what security requirements should be posed on individual, connected institutions? Are there e.g. concepts that could be adopted from the established European cooperation on passports and border control, to form a "Schengen network" for the European health sector?

The Norwegian Sector Norm [SHdir06] introduces issues related to admission policy. It establishes a framework of conformity in the security area that makes it possible for Norwegian health institutions to “lower their guard” and implement less stringent security measures against other institutions that also adhere to the norm. Enforcing the Sector Norm on all institutions connected to Norsk Helsenett enables the same “relaxed” security regime in the institutions’ health network connection, e.g. by permitting a single security barrier (firewall) instead of two barriers in a serial arrangement, as is required when connecting to other external networks.

In this perspective an ambition to create a Scandinavian health network with a minimum number of security barriers between any pair of health institutions may prove to be difficult to implement. The formal trust regime created by the Sector Norm may not be easily extendable across borders to other national health networks.

6.2 Infrastructure issues

As pointed out in section 2.1.2, use of NAT has a number of negative consequences. The requirement to use public IP addresses across the Danish S-DIX, and the policy to route only public IP addresses in Sjunet in reality mandates NAT for most of the data traffic that is carried in these networks. But the situation on Norwegian side of the border is also far from ideal. The end-to-end intention of using private IP addresses in a coordinated regime is being severely watered out by address collisions requiring “NAT islands”, both due to municipal health units under a local IP addressing regime connecting to the health network, and by inter-regional address conflicts dating back to the creation of Norsk Helsenett. In summary, the current state of affairs with respect to IP address use in the union of the national Scandinavian health networks is clearly not ideal.

The obvious solution is to acquire a sufficient number of public IPv4 addresses to cover all the hosts and networks involved, and to reconfigure all devices with these new, public addresses, and abolish NAT. But in light of the current shortage of IP addresses, and RIPE’s restrictive policy in handing out new addresses, this is hardly a realistic strategy. In fact, the shortage of IPv4 addresses is becoming so grave that APNIC has taken an initiative to establish a common procedure for all RIRs to terminate their allocation in a timely manner [Maemura07].

When this procedure is established, the argument for planning a conversion process to IPv6 addresses will become even more convincing. With the creation of a registration arrangement to ensure uniqueness of local IPv6 unicast addresses (ULA-C) [Martinez07], it is even possible to maintain the single desirable feature of the private IPv4 addresses: that they are not routed on the public Internet. There is no denying that a renumbering on a scale like this is a gigantic undertaking. Infrastructure must be reconfigured, and client systems need to be updated and possible replaced. Also, there are also a large number of server applications that have been taken into production use in the health sector over the years. The effort required to scrutinize and correct these has a potential to dwarf the y2k effort that provided good income for many consulting businesses over several years.

Equally significant and related to the IP addressing issue are the problems concerning DNS. Use of private IP addresses and NAT mandates a split horizon DNS, in which service names resolve into different IP addresses depending on where the question originates. In addition, the lack of a common DNS root in Norsk Helsenett makes service integration on the DNS level challenging.

It is difficult to see any other solution to the first of these issues than a migration to IPv6 addresses, possibly employing ULA-C. This would also enable a full integration between the health networks' internal DNSs and the Internet DNS, and making use of the inherent scaling mechanisms in the Internet DNS. This in turn would make irrelevant the DNS scaling issues related to a lack of a common root in Norsk Helsenett's DNS. The arguments left for converting to a common root are (1) the ability to easily recognize a domain name as health related, called the characterizing property of a common root in section 2.2, and (2) the possible vanity advantage of a having a ”topical” top level domain name like .health. However, if two-letter IS0-3166 country codes were used at the second level, then Norwegian health institutions would have their network appearances as sub-domains of no.health – clearly a less attractive option.

In fact, converting to IPv6 ULA-C addresses and merging all internal DNS data into the Internet DNS will have such positive effects on operational issues in Norsk Helsenett, that the effort may be justified even without the prospect of a Scandinavian health network cooperation.

In section 4.2 a categorization was made of the requirements that telemedicine applications present to a network, as message oriented, interactive or real-time. However, an application will not always have intrinsic communication requirements that unambiguously fit into one of these categories, but may depend on the user interface design philosophy of the application developer/programmer, or have properties that utilizes several categories. E.g. one can envisage a teleradiology application where interactivity is required for RIS access, but transfer of x-rays is done through EDI messages.

One of the fundamental questions is: should the Scandinavian health network provide real-time communication services to telemedicine applications that include a videoconferencing component, and therefore require service guarantees from the network? Credible delivery of such services requires implementation of QoS for selected traffic in each of the national networks, as well as IP level peering mechanisms between the national network service providers that honour these QoS mechanisms for cross-border traffic.

Telemedicine applications that are partly or wholly based on video conferencing abound. In fact, one out of four applications presented at the recent conference “Cross-border eHealth in the Baltic Sea Region” in Stockholm (www.ehealthconference.info) was based on videoconferencing, although transported by a network without QoS guarantees. The wide range of telemedicine applications where video conferencing is essential – meetings, remote education, medical consultations/second opinion, competence support during surgical procedures, telepresence aid for geographically dispersed groups, and more – caters for the requirement of any modern health network to include service support for real-time traffic as an integrated part of the infrastructure.

However, considering the limited amount of experience with QoS provisioning in the national health networks in the Scandinavian countries, it may not be realistic to put this item on the short-term agenda for a Scandinavian health network. Also, the specification of QoS mechanisms between peering network providers has been a research topic for several years, although the mechanisms necessary to accomplish such arrangements are just now beginning to surface from the standardization processes. It may seem that the path towards QoS guaranteed network services in a Scandinavian health network is some years into the future.

6.3 Motivational issues

At the end of chapter 2, examples were provided of existing and passed cross-border e-health activities. At the bottom line, the motivation for extending health networks across country borders will have to rise from the amount of useful applications that can be carried, and the benefits that they bring. But to envisage and describe potential applications is no easy task. Besides, there is also a chicken-and-egg aspect present: building a cross-border network without a sufficiently convincing application portfolio vs. imagining a "killer application" without a network available to run it over.

One category of applications that is bound to become important in cross-border networking is also exemplified in section 2.6: purchase of diagnostic services. The example describes how TMC in Barcelona offers its radiology expertise over distance, but in fact Namsos/Levanger had a similar arrangement with specialists in Sweden simultaneously. A wide range of telemedicine services could be made subject to trade on an international health market, as they many places have been for some time on a regional or national basis, motivated by lack of qualified personnel in small hospitals, 24-by-7 emergency service cooperation etc.

There is also another aspect of the diagnostic service purchase model that extends beyond a mass marked. As in many other sectors, the trend towards specialization is common in the health area, creating narrow pockets of highly specialized skills in centres that operate in a global health market. These centres treat rare cases and need large uptake areas, some of which may cross borders. Access to this competence will become increasingly important in the process of providing steadily improved health services to the population.

6.4 Further work

Although it proved impossible to execute the performance measurements as planned, there is still a general consensus among the involved parties that it would still be interesting to conduct these measurements over the S-DIX based network. Requests for firewall reconfiguration are still pending, and more relevant and valid performance measurements will be conducted when the resource situation improves. In the mean time, the experienced TCP window and UDP problems on the IxChariot platform should be sorted out.

This thesis work only deals with a narrow set of technical aspects involved in building a Scandinavian health network. But the discussion on how these issues should be resolved, should be undertaken in a broader forum. Also, many important technical issues like carrier technology, QoS regime and security solutions are barely mentioned in the present work, and need considerable more surveying and discussion.

Other related issues are on the agenda in other contexts. Some legal issues concerning cross-border health cooperation are identified and discussed as part of the Baltic eHealth project in [Nohr05], as are organizational and financial issues [Linstad07]. And several groups and institutions (including MedCom) are exploring practical aspects of medical cooperation across national borders, as was demonstrated at the ehealthconference.info conference. But there may still be issues concerning the formal status and political anchoring of a cooperative effort to create a Scandinavian health network that need additional focus.

7 Concluding remarks

There are good reasons for commencing a process of harmonization of the infrastructure in the Scandinavian health networks. Employing IPv6, possibly with ULA-C addresses, and a DNS fully integrated with the Internet DNS are two elements that appear to be necessary for a successful and flexible interconnection between the networks. This will re-establish the end-to-end principle in the networks, and provide a flexible technological platform for further expansion to other regional or national health networks.

The costs associated with such an address migration will obviously be very high. But there is also the question: can the cost be avoided? Is the question really *when* a network IP version change should be installed, rather than *if*? Furthermore, there are strong indications that the answer to the when-part of the question is *soon* [Potaroo07].

Regardless of the outcome of such a harmonization process, closer contact between the health network organizations should be considered. The Scandinavian countries have had mutual benefits of cooperation in a number of areas. Even if the creation of an integrated Scandinavian health network should prove to be a premature idea, all parties are likely to benefit from establishing closer contact – possibly through an network operations and management forum.

On a more general note, the question might be raised: why have 100% of the Scandinavian countries established national health networks, when the idea has gained so little popularity in the rest of the world? Apart from the traditional explanations of well-organized Scandinavian societies, cultural similarities and a tradition of looking to each other for good ideas – are there aspects of the national health network concept that makes it less suitable in other parts of the world? There may be “food-for-thought” for sociologists here.

References

[Abdirad06] Afsin Abdirad, Babak Sarrafpour, Siavash Ghaderi-sohi. Static telepathology in cancer institute of Teheran University: report of the first academic experience in Iran. Digital Pathology 2006, 1:33. http://www.diagnosticpathology.org/content/pdf/1746-1596-1-33.pdf, accessed May 2007.

[Abelsen06] Lisbeth Remlo Abelsen, Arnstein Vestad, Daniel Nygård, Eva Skipenes, Harald Øverli Eriksen, Leif Erik Nohr, Line Nordgård. Ses@m Tromsø – eventyrlige muligheter for pleie- og omsorgstjenesten? Nasjonalt senter for telemedisin, 2006.

[AG-DHA06] Australian Government, Department of Health and Ageing. The Broadband for Health Programme - Overview. http://www.health.gov.au/internet/wcms/publishing.nsf/Content/health-ehealth-broadband-initiative.htm, accessed May 2007.

[Albitz01] Paul Albitz, Cricket Liu. DNS and BIND, 4th edition. O'Reilly & Associates, Inc, 2001. (ISBN: 0-596-00158-4).

[Anvari05] Mehran Anvari, Craig McKinley, Harvey Stein. Establishment of the World's First Telerobotic Remote Surgical Service: For Provision of Advanced Laparoscopic Surgery in a Rural Community. Annals of Surgery. 241(3):460-464, March 2005

[Arvidsson06] Lars Arvidsson, Björn Widenberg. Carelink: Genomlysning av Sjunet, version 1.3, 2006-04-03. http://www.carelink.se/dokument/forvaltning_och_tjanster/sjunet/doc_2006424130144.pdf, accessed April 2007.

[Baltic07] Baltic Health Network. http://www.baltic-ehealth.org/Baltic_Health_Network.htm, accessed April 2007.

[Bech02] Martin Bech, Ib Lucht. Teknisk forundesøgelse vedr. det Internetbaserede sundhedsdatanet Version 1.1. UNI-C april 2002. http://www.medcom.dk/dwn389, accessed April 2007.

[Bech03] Martin Bech, Ib lucht. Det internetbaserede sundhedsdatanet Vesion 1.0. UNI-C December 2003. http://www.medcom.dk/dwn390, accessed April 2007.

[Bech07] Martin Bech. The connection agreement system – in just 30 minutes. Presentation at the HDN.eu meeting 9th January 2007. http://hdn.eu/presentation/Denmark_Martin_HDN.eu_070109.ppt, accessed April 2007.

[Bergstrøm04] Roald Bergstrøm, project manager. Digital røntgen: Distribuert granskning av røntgenbilder over bredbånd. Norges Forskningsråd, Final report Høykom project no 1605/240. http://www.hoykom.no/hoykom/HOYKOM_Prosjekter_ny.nsf/%0B8a047f65e9984d87c1256d510048e307/c92a98159586c497c1256e2a00573a93/$FILE/P1605%20Digital%20Røntgen.doc, accessed May 2007.

[Calyam05] Prasad Calyam, Mukundan Sridharan, Weiping Mandrawa, Paul Schopis. Performance Measurements and Analysis of H.323 Traffic. Lecture Notes in Computer Science, Vol. 3015, Springer, 2004 (ISBN 978-3-540-21492-2).

http://www.osc.edu/research/networking/PDFs/h323.pdf, accessed May 2005.

[Carelink03] Carelink. Sjunet Riktlinjer för säkerhet. CIS 1/2003. http://www.carelink.se/dokument/forvaltning_och_tjanster/sjunet/doc_2003317083959.pdf, accessed April 2003.

[Carelink04] Carelink. Slutrapport upphandling nytt nät. Sjunet Datakom CIS 8/2004. http://www.carelink.se/dokument/forvaltning_och_tjanster/sjunet/doc_2005110154358.pdf, accessed April 2004.

[Carelink05a] Carelink. Tekniska anvisningar nyttjare. Sjunet Datakom Vers. 2.1, 2005-07-06. http://www.carelink.se/dokument/forvaltning_och_tjanster/sjunet/doc_200575160449.pdf, accessed April 2007.

[Carelinke05b] Carelink. Sjunet – Nytta anslutning. Version 1.6 2005-06-29. http://www.carelink.se/dokument/forvaltning_och_tjanster/sjunet/doc_2005629161019.pdf, accessed April 2007.

[Carelink07a] CareLink. Bastjänst HSA. http://www.carelink.se/tjanster/hsa/, accessed March 2007.

[Carelink07b] Carelink. Tjänster – Förvaltningsområden. http://www.carelink.se/tjanster/, accessed April 2007.

[Carelink07c] Carelink. Företag och organistationer som erbjuder tjänster på Sjunet. http://www.carelink.se/tjanster/leverantorstjanster/, accessed April 2007.

[Carelink07d] Carelink. Medlemmar. http://www.carelink.se/om_carelink/medlemmar/, accessed April 2007.

[CDNNCC07] National Collaborating Centre, Determinants for Health. About the NCC's. http://www.nccdh.stfx.ca/about-nccs.htm, accessed May 2007.

[CFST07] Center for Sundhedstelematik. MedCom – det danske sundhedsdatanet. http://www.cfst.dk/wm143353, accessed January 2007.

[Chafe00] Chris Chafe, Scott Wilson, Randal Leistikow, Dave Chisholm, Gary Scavone. A simplified approach to high quality music and sound over IP. Proceedings of the COST-G6 Conference on Digital Audio Effects (DAFX-00), Verona, Italy, December 7-9, 2000. http://www-ccrma.stanford.edu/~cc/misc-papers/dafx2000.ps, accessed April 2007.

[Chall03] B.J. Challacombe, L.R. Kavoussi, P. Dasgupta (2003). Trans-oceanic telerobotic surgery. BJU International 92 (7), 678-680

[Cisco01] Cisco VPN Solutions. Cisco Systems Inc. 2001. http://www.cos-cug.org/Presentations/VPN_Site-2-Site.ppt, accessed April 2007.

[Conn#5] Connected #5. NHS N3 Mailing List Bulletin. http://www.n3.nhs.uk/files/bulletins/connectedissue05.pdf, accessed May 2007.

[Cottrell02] Les Cottrell. Monitoring Internet connectivity of Research and Educational Institutions. http://www.ictp.trieste.it/~ejds/seminars2002/Les_Cottrell/ictp-02-final.ppt, accessed March 2007.

[DaVinci05] The da Vinci Surgical System.

http://www.intuitivesurgical.com/products/davinci_surgicalsystem/index.aspx, accessed April 2007.

[Debian06] Debian Bug report logs - #353037; iperf: does not work in daemon mode. *http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=353037*, accessed May 2007.

[DGINFSO06] EU Information Society and Media Directorate-General. Sollefteå and Borås hospitals; Sjunet, Sweden – radiologyconsultations between Sweden and Spain. eHealth impact 7.10, DG INFSO October 2006. *http://ec.europa.eu/information_society/activities/health/docs/events/opendays2006/ehealth-impact-7-10.pdf*, accessed April 2007.

[Divaharan05] Diva Divaharan. N3 Service Description (end user) – IP Address Allocation Process, Issue 1.4. British Telecommunications Plc, 30 November 2005. *http://n3.nhs.uk/files/technicalguidance/N3IPAddressallocationprocess.doc*, accessed March 2007.

[Eriksson07] Johnny Eriksson. Telemedicine "plug-in" en förlängning av den egna verksamheten. Presentation at Röntgenveckan 2006, Örebro. *http://www.rontgenveckan.se/2006/downloads/telmedicine%20plugin.pdf*, accessed May 2007.

[EUComm04] Commission of The European Communities. Communication from The Commission to The Council, The European Parliament, The European Economic And Social Committee and The Committee Of The Regions; e-Health – making healthcare better for European citizens: An action plan for a European e-Health Area. COM(2004) 356 final, 30.4.2004. *http://europa.eu.int/eur-lex/lex/LexUriServ/site/en/com/2004/com2004_0356en01.pdf*, accessed May 2004.

[Fitchett05] Trevor Fitchett, John Potts. QoS: The N3 Framework. *http://www.n3.nhs.uk/n3scotland/files/generalfiles/N3_Scotland_the_QoS_Framework.ppt*, accessed May 2007.

[FreeBSD02] FreeBSD Hypertext Man Pages. dummynet(4), October 28, 2002. *http://www.freebsd.org/cgi/man.cgi?query=dummynet&sektion=4*, accessed May 2007.

[FreeBSD06] FreeBSD Hypertext Man Pages. ipfw(8), July 25, 2006. *http://www.freebsd.org/cgi/man.cgi?query=ipfw&sektion=8&apropos=0&manpath=FreeBSD+6.2-RELEASE*, accessed May 2007.

[Gamst07] Oddvar Gamst, UNN. Personal communication.

[Glenholmes07]Tony Glenholmes, NHS. Personal communication.

[Haglund07] Sten Haglund, Landstinget Västmanland. Personal communication.

[Haug97] Steinar Haug, Olaf Schjelderup. Noen tekniske og sikkerhetsmessige aspekter vedrørende realisering av et nasjonalt helsenett. 1997-10-30. UNINETT project no. 359501.00.

[Hausken91] Peter Hausken. TelemaX.400 settes i offisiell drift. Uninytt 1-1991. *http://forskningsnett.uninett.no/uninytt/1991-1/tel.html*, accessed April 2007.

[HCFAQ07] HealthConnect. FAQs.

http://www.health.gov.au/internet/hconnect/publishing.nsf/Content/faqs-1lp, accessed May 2007.

[HDN07] HDN.eu Documents. *http://hdn.eu/documents/*, accessed April 2007.

[HealthLink02] HealthLink Ltd. Company Profile, September 2002. *http://www.healthlink.net/healthlink_documents/brochure/HealthLink%20Profile.zip*, accessed May 2007.

[Hinden05] R. Hinden, B. Haberman. Centrally Assigned Unique Local Ipv6 Unicast Addresses, February 2005. Work in progress. *http://tools.ietf.org/html/draft-ietf-ipv6-ula-central*, accessed May 2007.

[HISAC06] Health Network. About the Health Network, 24/08/2006. *http://www.hisac.govt.nz/moh.nsf/indexcm/hisac-network-home*, accessed May 2007.

[HoD96] Helse- og omsorgsdepartementet. Mer helse for hver bIT, Informasjonsteknologi for en bedre helsetjeneste. *http://www.regjeringen.no/nb/dep/hod/dok/Veiledninger_og_brosjyrer/1996/Mer-helse-for-hver-bIT.html?id=87401*, accessed March 2007.

[Hætta07] Rune Hætta, Norsk Helsenett AS. Personal communication.

[Høykom03] Høykom, Norges forskningsråd. Kompetanseheving i fosterdiagnostikk ved hjelp av telemedisin. Final report Høykom project no 1372/240. *http://www.hoykom.no/hoykom/HOYKOM_Prosjekter_ny.nsf/b1d3714b8cf10cedc1256d51004995c2/c02e9b390f66159041256c0e003e58f7/$FILE/P1372%20Sluttrapport%20ultranett.doc*, accessed April 2007.

[Ixia06] IxChariot User Guide, Release 6.40.

[Ixia07] IxChariot. *http://www.ixiacom.com/products/ixchariot/*, accessed April 2007.

[ITU00] One-way Transmission Time. ITU-T Rec. G. 114, ITU Geneva, May 2000.

[ITU06a] IP packet transfer and availability performance parameters. ITU-T Rec. Y. 1540, ITU Geneva, July 2006.

[ITU06b] Networl Performance Objectives for IP-Based Services. ITU-T Rec. Y. 1541, ITU Geneva, July 2006.

[Krogsrud07| Hans Petter Krogsrud, Norsk Helsenett AS. Personal communication.

[LeParc02] Philippe Le Parc, Pascal Ogor, Jean Vareille, Lionel Marcé. Web Based Remote Control of Mechanical Systems. IEEE 2002 International Conference on Software, Telecommunications and Computer Networks (SofCOM'02). *http://www.ea2215.univ-brest.fr/publications/par_chercheur/le_parc/le_parc_Softcom2002.pdf*, accessed May 2007.

[Linstad07] Line Linstad, Elin Breivik. Organisational and financial challenges in eHealth services. Presented at *Cross-border eHealth in the Baltic Sea Region*, Stockholm 21-22 May, 2007. *http://www.ehealthconference.info/Presentations/c_line_linstad.pdf*, accessed May 2007.

[Lother05] Ann Therese Lotherington (red). Telemedisin i pleie- og omsorgstjenesten: Om å takle det uforutsette; Rapport fra prosjektet SES@m Tromsø. NORUT Samfunnsforsking AS / Nasjonalt senter for telemedisin, Report no 11/2005. *http://www.telemed.no/getfile.php/223000.357.payxudbsxd/SES%40m+Troms%F8+-+Midtveisrapport.pdf*, accessed April 2007.

[Maemura07] Akinori Maemura, Toshiyuki Hosaka, Takashi Arano, Kuniaki Kondo, Tomohiro Fujisaki, Kosuke Ito, Shuji Nakamura, Tomoya Yoshida, Susumu Sato, Akira Nakagawa. IPv4 Countdown Policy. RIPE Policy Proposal 2007-03, 24 April 2007. *http://www.ripe.net/ripe/policies/proposals/2007-03.html*, accessed April 2007.

[Malmqvist04] Gustav Malmqvist, K.G. Nerander, Mats Larson. Sjunet – The National IT Infrastructure for Healthcare in Sweden. Studies in Health Technology and Informatics, Vol. 100, IOS Press, 2004 (ISBN: 1 58603 448 0).

[Martinez07] Jordi Palet Martinez. IPv6 ULA-Central. RIPE Policy Proposal 2007-05, 2 May 2007. *http://www.ripe.net/ripe/policies/proposals/2007-05.html*, accessed May 2007.

[MedCom01] MedCom . Fremtidens sundhedskommunikation. 2001. *http://www.medcom.dk/dwn149*, accessed April 2007.

[MedCom04] MedCom . Introduktion til det internetbaserede Sundhedsdatanet. 2004. *http://www.medcom.dk/dwn184*, accessed April 2007.

[MedCom06a]MedCom-status i antal. *http://www.medcom.dk/wm110164*, accessed March 2007.

[MedCom06b] MedCom. Sundheds-DIX. *http://www.medcom.dk/default.asp?id=110002*, accessed December 2006.

[MedCom07a]MedCom. Om os. *http://www.medcom.dk/wm109974*, accessed January 2007

[MedCom07b] MEDCOM løsningsforslag til videokonference hos Vejle Amt. *http://www.medcom.dk/dwn557*, accessed April 2007.

[MedCom07c]MedCom 5, Sundheds-DIX. *http://www.medcom.dk/wm110045*, 04-04-2007, accessed April 2007.

[MedCom07d] MedCom. Katalog over servere på SDN "SundhedsDIXen". March 2007. *http://www.medcom.dk/dwn711*, accessed April 2007.

[MedCom07e]MedCom. Fejlfinding og support. http://www.medcom.dk/wm110405 , accessed May 2007.

[MedCom07f] MedCom. Internationale projekter. *http://www.medcom.dk/wm109986*, accessed May 2007.

[Mogens07] Kasper Mogensen, Ib Lucht. Aftalesystem for sundhedsdatanettet Version 3.0. UNI-C March 2007. *http://www.medcom.dk/dwn611*, accessed April 2007.

[NCM05] Health and Social Sectors with an "e"; A study of the Nordic countries. Nordic Council of Ministers, Copenhagen 2005 (ISBN 92-893-1157-6). *http://www.carelink.se/dokument/internationellt/doc_2005412151704.pdf*, accessed April 2007.

[NHN07a] Norsk Helsenett. *http://www.nhn.no/*, accessed March 2007.

[NHN07b] Norsk Helsenett. Adresseregisteret i Norsk Helsenett. http://www.nhn.no/Tjenester/adresseregister/index_html, accessed March 2007.

[NHNWAN05] Kravspesifikasjon WAN Norsk Helsenett. Fjerde utkast 12. august 2005.

[NLANR05] NLANR Distributed Applications Support Team, Iperf project page. http://dast.nlanr.net/Projects/Iperf/, accessed May 2007.

[Nohr05] Leif Erik Nohr, Manolis Nymark, Marika Zmenja. Report on identified legal issues of the Baltic eHealth project. http://www.baltic-ehealth.org/news/Publications/Baltic_eHealth_Legal_issues_Report.pdf, accessed March 2007.

[Nohr06] Leif Erik Nohr, Adam Martony, Line Linstad, Elin Breivik, Ernst Kloosterman. Cross border eHealth in the Baltic Sea Region – what issues should be considered?, June 26, 2006. http://www.baltic-ehealth.org/intern/wp1/Guidelines/Cross_border_eHealth_in_the_Baltic_Sea%20_Region_report_june2006.pdf, accessed May 2007.

[Pedersen05] Claus Duedal Pedersen. An Baltic healthcare network and interoperability challenges. Presented at *Cisco eHealth think tank meeting*, January 2005. http://www.baltic-ehealth.org/news/Ehealth-thinktank_25_01_05.ppt, accessed March 2007.

[Pedersen07] Claus Duedal Pedersen, MedCom. Personal communication.

[PHAoC07] Public Health Agency of Canada. Health Portfolio. http://www.phac-aspc.gc.ca/portfolio_e.html, accessed May 2007.

[Phifer00] Lisa Phifer. The Trouble with NAT. The Internet Protocol Journal No. 3 Vol. 4 December 2000. http://www.cisco.com/web/about/ac123/ac147/ac174/ac182/about_cisco_ipj_archive_article09186a00800c83ec.html, accessed May 2007.

[PITAC04] President's Information Technology Advisory Committee. Report to The President, Revolutionizing Health Care Through Information Technology. 2004. http://www.nitrd.gov/pitac/reports/20040721_hit_report.pdf, accessed May 2007.

[Potaroo07] IPv4 Address Report. http://www.potaroo.net/tools/ipv4/, accessed May 2007.

[Ramel06] Björn Ramel. Robotkirurgin ökar – men utvärderinga saknas. Läkartidningen, Organ för Sveriges läkarförbund, page 3788 06-11-29. http://www.lakartidningen.se/engine.php?articleId=5591, accessed April 2007.

[Ranum99] Marcus J. Ranum. Problems with the Firewall model. http://www.ranum.com/security/computer_security/archives/problems-with-firewalls.pdf, accessed April 2007.

[RFC768] J. Postel. User Datagram Protocol. August 1980. http://tools.ietf.org/html/rfc0768, accessed March 2007.

[RFC791] J. Postel. Internet Protocol. September 1981.

http://tools.ietf.org/html/rfc0791, accessed March 2007.

[RFC793] J. Postel. Transmission Control Protocol. September 1981. http://tools.ietf.org/html/rfc0793, accessed April 2007.

[RFC805] J. Postel . Computer mail meeting notes. February 1982. http://tools.ietf.org/html/rfc0805, accessed May 2007.

[RFC1034] P.V. Mockapetris. Domain names - concepts and facilities. November 1987. http://tools.ietf.org/html/rfc1034, accessed February 2007.

[RFC1035] P.V. Mockapetris. Domain names - implementation and specification. November 1987. http://tools.ietf.org/html/rfc1035, accessed February 2007.

[RFC1323] V. Jacobson, R. Braden, D. Borman. TCP Extensions for High Performance. http://tools.ietf.org/html/rfc1323, accessed May 2007.

[RFC1631] K. Egevang, P. Francis. The IP Network Address Translator (NAT). May 1994. http://tools.ietf.org/html/rfc1631, accessed March 2007.

[RFC1633] R. Braden, D. Clark, S. Shenker. Integrated Services in the Internet Architecture: an Overview. June 1994. . http://tools.ietf.org/html/rfc1633, accessed May 2007.

[RFC1889] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson. RTP: A Transport Protocol for Real-Time Applications. Audio-Video Transport Working Group, January 1996. http://tools.ietf.org/html/rfc1889, accessed March 2007.

[RFC1900] B. Carpenter, Y. Rekhter. Renumbering Needs Work. February 1996. http://tools.ietf.org/html/rfc1900, accessed May 2007.

[RFC1918] Y. Rekhter B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear. Address Allocation for Private Internets. February 1996. http://tools.ietf.org/html/rfc1918, accessed March 2007.

[RFC1958] B. Carpenter, ed. Architectural Principles of the Internet. June 1996. http://tools.ietf.org/html/rfc1958, accessed May 2007.

[RFC2101] B. Carpenter, J. Crowcroft, Y. Rekhter. IPv4 Address Behaviour Today. February 1997. http://tools.ietf.org/html/rfc2101, accessed March 2007.

[RFC2205] R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, S. Jamin. Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification. September 1997. http://tools.ietf.org/html/rfc2205, accessed May 2007.

[RFC2330] V. Paxson, G. Almes, J. Mahdavi, M. Mathis. Framework for IP Performance Metrics. May 1998. http://tools.ietf.org/html/rfc2330, accessed April 2007.

[RFC2475] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss. An Architecture for Differentiated Service. December 1998. http://tools.ietf.org/html/rfc2475, accessed May 2007.

[RFC2663] P. Srisuresh, M. Holdrege. IP Network Address Translator (NAT) Terminology and Considerations. August 1999. http://tools.ietf.org/html/rfc2663, accessed February 2007.

[RFC2680] G. Almes, S. Kalidindi, M. Zekauskas A One-way Packet Loss Metric for IPPM, September 1999. http://tools.ietf.org/html/rfc2680, accessed May 2007.

[RFC2681] G. Almes, S. Kalidindi, M. Zekauskas. A Round-trip Delay Metric for IPPM. September 1999. http://tools.ietf.org/html/rfc2681, accessed May 2007.

[RFC2775] B. Carpenter. Internet Transparency. February 2000. http://tools.ietf.org/html/rfc2775, accessed April 2007.

[RFC2993] T. Hain. Architectural Implications of NAT. November 2000. http://tools.ietf.org/html/rfc2993, accessed February 2007.

[RFC2998] Y. Bernet, P. Ford, R. Yavatkar, F. Baker, L. Zhang, M. Speer, R. Braden, B. Davie, J. Wroclawski, E. Felstaine. A Framework for Integrated Services Operation over Diffserv Networks. November 2000. http://tools.ietf.org/html/rfc2998, accessed May 2007.

[RFC3022] P. Srisuresh, K. Egevang. Traditional IP Network Address Translator (Traditional NAT). January 2001. . http://tools.ietf.org/html/rfc3022, accessed February 2007.

[RFC3027] M. Holdrege, P. Srisuresh. Protocol Complications with the IP Network Address Translator. January 2001. http://tools.ietf.org/html/rfc3027, accessed March 2007.

[RFC3102] M. Borella, J. Lo, D. Grabelsky, G. Montenegro. Realm Specific IP: Framework. October 2001. http://tools.ietf.org/html/rfc3102, accessed May 2005.

[RFC3724] J. Kempf, R. Austein, eds., IAB. The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture. March 2004. http://tools.ietf.org/html/rfc3724, accessed April 2007.

[RFC3393] C. Demichelis, P. Chimento. IP Packet Delay Variation Metric for IP Performance Metrics (IPPM). November 2002. . http://tools.ietf.org/html/rfc3393, accessed May 2007.

[RFC4193] R. Hinden, B. Haberman. Unique Local IPv6 Unicast Addresses. October 2005. http://tools.ietf.org/html/rfc4193, accessed May 2007.

[RIPE-388] IPv6 Address Allocation and Assignment Policy. RIPE, September 2006. ftp://ftp.ripe.net/ripe/docs/ripe-388.pdf, accessed May 2007.

[RIPE-405] IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region. RIPE, April 2007. ftp://ftp.ripe.net/ripe/docs/ripe-405.pdf, accessed May 2007.

[Saltzer84] J.H. Saltzer, D.P. Reed, D.D. Clark. End-To-End Arguments in System Design. ACM TOCS Vol. 2 No. 4, November 1984. http://www.reed.com/Papers/EndtoEnd.html, accessed May 2007.

[SHdir03] Nasjonalt helsenett; Konkurransegrunnlag. Sosial- og helsedirektoratet 03.02.2003.

[SHdir06] Norm for informasjonssikkerhet i helsesektoren. Sosial- og helsedirektoratet. August, 2006. http://www.nhn.no/Tjenester/bransjenormen/filer/norm_for_informasjonssikkerhet_i_helsesektoren_7august2006.pdf, accessed March 2007.

[Siticom02] Etablering af national sundhedsportal. Projektbeskrivelse Version 1.0 1. februar 2002. Siticom group.

http://www.sundhed.dk/Images/alle/redaktion/pdf/sundhedsportal.pdf, accessed April 2007.

[Skjetne07] Atle Skjetne, Helse Nord-Trøndelag HF. Personal communication.

[Slot07] Dr. Marjan Slot, Eurotransplant Medical Staff. Personal communication.

[SOIL07] The SOIL Network. http://www.oilcamp.com/portal/Portals/0/docs/SOIL%20Datasheet.pdf, accessed March 2007.

[Sorth03] Lennart Sorth. Tilslutning til Sundhedsdatanettet (SDN). Notat 13.08.03. http://www.medcom.dk/dwn610, accessed April 2007.

[Sørensen04] Frode Sørensen. Moderne IP-nett. IDG Norge Books AS, 2004. ISBN-82-7772-279-6.

[Sundhed05] Sundhedsstyrelsen. SOR – Sundhedsvæsenets Organisationsregister 2005. Version 1.0 July 5. 2005 (ISBN 87-7676-133-9) http://www.sundhedsstyrelsen.dk/upload/informatik_og_sundhedsdata/sundhedsinformatik/klassifikationer/sorv1.pdf, accessed May 2007.

[Uninett06] Uninett. Målepåle – måleinfrastruktur, 2006-04-18. http://forskningsnett.uninett.no/produkt/maalepale.html, accessed March 2007.

[Wack02] John Wack, Ken Cutler, Jamie Pole. Guidelines on Firewalls and Firewall Policy. National Institute of Standards and Technology. Special Publication 800-41, January 2002. http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf, accessed May 2007.

[Wanscher07] Christina E. Wanscher and Henning Voss. The Baltic eHealth Project; Connecting Regions - Optimising healthcare in the Baltic Sea Region. Hospital - The Official Journal of the European Association of Hospital Managers 4/2005. http://www.baltic-ehealth.org/news/press/H4_IT3_OE_Baltic%20eHealth.pdf, accessed May 2007.

[WHO06] World Health Organization. Building Foundations for eHealth; Progress of member states. © World Health Organization 2006 (ISBN 978-92-4-159504-9). http://www.who.int/ehealth/resources/bf_full.pdf, accessed May 2007.

[XIWT98] Cross-Industry Working Team. Customer View of Internet Service Performance: Measurement Methodology and Metrics, October 1998. http://www.xiwt.org/documents/IPERF-paper.pdf, accessed March 2007.

[Yorine07] Yorine Internet Information Center, AS Number 224. http://yorine.nl/network/as/224, accessed May 2007.