

CONTRACT NUMBER 031513

eDEISA
 EXTENDED DISTRIBUTED EUROPEAN
 INFRASTRUCTURE FOR
 SUPERCOMPUTING APPLICATIONS

European Community Sixth Framework Programme
 RESEARCH INFRASTRUCTURES
 Integrated Infrastructure Initiative

eSA1: Second 10Gb/s network infrastructure report

Deliverable ID: eDEISA_D-eSA1-2

Due date : May, 31st, 2007
Actual delivery date: May, May 31st, 2007
Lead contractor for this deliverable: FZJ, Germany

Project start date : June 1st, 2006
Duration: 2 years

| | | |
|--|---|---|
| Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006) | | |
| Dissemination Level | | |
| PU | Public | |
| PP | Restricted to other programme participants (including the Commission Services) | X |
| RE | Restricted to a group specified by the consortium (including the Commission Services) | |
| CO | Confidential, only for members of the consortium (including the Commission Services) | |

Table of Contents

| | |
|--|-----------|
| Table of Contents..... | 1 |
| 1 INTRODUCTION | 2 |
| 2 THE ESA1 – EXTENDED NETWORK INFRASTRUCTURE DESIGN | 2 |
| 3 MANAGING A 10 GB/S EUROPEAN NETWORK INFRASTRUCTURE..... | 5 |
| 4 OPERATION EXPERIENCES, THROUGHPUT VALUES, ET AL..... | 6 |
| 5 NETWORK DESIGN AND ITS INFLUENCE ON APPLICATIONS..... | 7 |
| 6 LESSONS LEARNED | 11 |
| 7 FUTURE ENHANCEMENTS AND CO-OPERATIONS | 12 |
| 8 SUMMARY | 13 |
| 9 REFERENCES AND APPLICABLE DOCUMENTS..... | 13 |
| 10 LIST OF ACRONYMS AND ABBREVIATIONS | 13 |

1 Introduction

The fundamental objective of the eSA1 Service Activity is to enable the deployment of a multi-gigabit per second network infrastructure for the DEISA supercomputing Grid [1].

The DEISA SA1 activity compiled detailed specifications for the dedicated bandwidth network infrastructure needed for the operation of the DEISA supercomputing Grid in phase 1 and phase 2 of the DEISA project. These specifications were evaluated and implemented together with the National Research Networks (NRENs) directly concerned and with DANTE, the operational interface of GEANT/GN2 [2][3]. Within the DEISA SA1 activity a second “proof of concept” phase connecting the five DEISA sites FZJ, IDRIS, LRZ, RZG and SARA was initiated which paved the way to an extension of 10 Gbit/s connectivity to all the other DEISA sites.

All the activities within DEISA and eDEISA are dependent on the technical and contractual availability of fiber links and especially wavelength division equipment.

Due to delays in the delivery of equipment, the second DEISA “proof of concept” phase could not be started on time. A central DEISA switch has been installed in Frankfurt and first tests have been made in November 2006. In the meantime seven sites could be connected to the 10 Gbit/s DEISA backbone, namely BSC, FZJ, HLRS, IDRIS, LRZ, RZG, and SARA, and have gone into production without any problems.

This report describes the activities done so far in eSA1, the current layout of the DEISA backbone, and the future plans for infrastructure enhancements.

2 The eSA1 – Extended Network Infrastructure design

At start of the DEISA project in 2004 it has been very expensive to install a 10 Gb/s network infrastructure across Europe. NRENs and GÉANT have not been capable of technically providing cheap links across national boundaries. Therefore the DEISA consortium decided to start with a 1 Gb/s network infrastructure using “virtually dedicated” links between DEISA sites. Being aware that a fully-meshed topology involving point-to-point connections across all platforms can hardly be considered, because of involved costs, GÉANT proposed the use of a PREMIUM IP service [4] provided via the normal GÉANT infrastructure. This design led to a full meshed 1 Gb/s network infrastructure between involved NRENs of participating DEISA sites. DEISA started the implementation as a “proof of concept” phase including only the sites CINECA, FZJ, IDRIS and RZG to deeply analyse the communication behaviour and to prepare a full mesh to all sites in future. A multilayer switch had been installed at every DEISA site to which the local supercomputer systems had been connected by one or more Gigabit-Ethernet interfaces (Etherchannel). With this design and using the Premium IP service an upper-bounded one-way delay, upper-bounded Instantaneous Packet Delay Variation (IPDV), no packet loss due to congestion and guaranteed capacity as well as a throughput capacity of 1 Gb/s could be guaranteed. After careful analysis the “proof of concept” phase could be completed successfully and a 1 Gb/s network infrastructure to the remaining sites has been established. The network phase 1 is now working for three years in a stable production status without larger disruptions and failures.

In July 2005 a first DEISA-GÉANT workshop was held in Munich to discuss a 10 Gb/s network infrastructure for DEISA. Here a small working group had been initiated which considered design, backup scenarios, and scenarios for potential future extensions concerning bandwidth and number of participants.

In the beginning of 2006 it became technically feasible to start negotiations with all the providers involved. The negotiations have been facilitated because of the availability of the new GÉANT2 infrastructure allowing multiple wavelengths across the installed GÉANT2 footprint. The DWDM design of the European infrastructure provided the ability to have additional wavelengths without major additional technical efforts. It became obvious to design the new DEISA network in a star like fashion with a central DEISA multilayer switch located somewhere central in Europe. Contracts could be signed with the German NREN DFN to provide a central placement of the DEISA-10GE-switch at a housing company at Frankfurt, Germany, where DFN- and GEANT2-PoPs were already established, because four of the DEISA partners are located in Germany. Edge layer 2/3 switches have been installed at all DEISA sites. A direct connectivity between these switches, that is wavelength across NRENs and GÉANT2 DWDM equipment, offers short delays, simple management and low total cost of ownership.

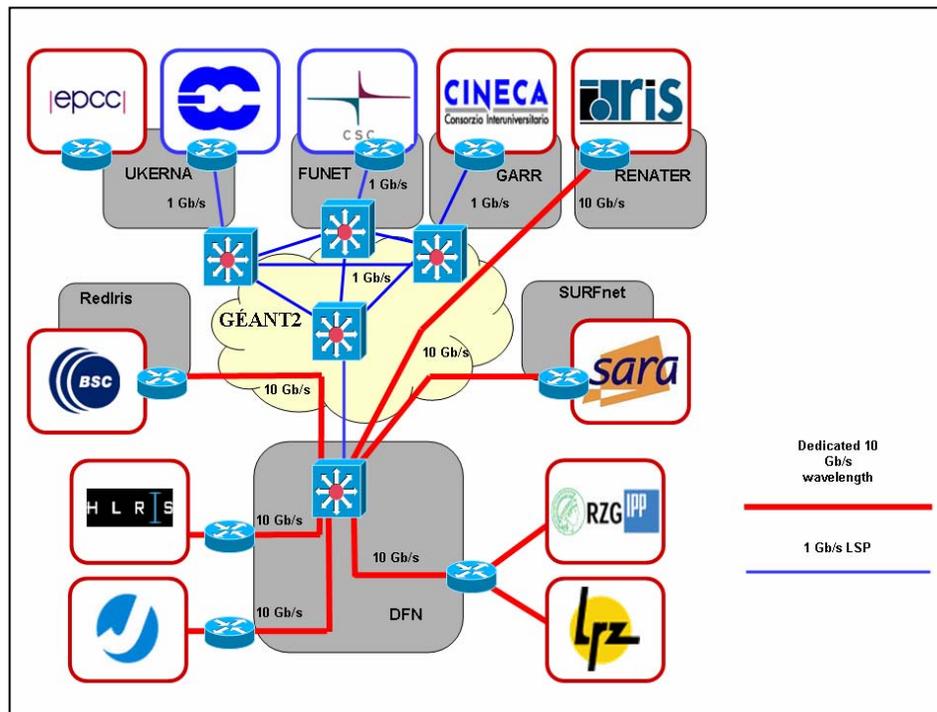


Figure 1: DEISA Network (technical overview April 2007)

A Cisco Catalyst 6500 switch has been chosen as heart of the DEISA network infrastructure. The installed network links are 10 Gb/s Ethernet. The 10GE-interfaces which have XENPACK-10GB-LR optics at Frankfurt have been configured with FlowControl and JumboFrames enabled. They do not issue nor accept BPDUs. The links between central switch and edges use private network addresses as defined in RFC 1918 [5]. An uplink to the worldwide INTERNET has been configured, to allow easy integration of the DEISA partner sites which do not have 10 Gb/s links yet. Access lists installed on every interface of the central switch allow a secure network environment. Any DEISA site is open to install additional firewall(s) or packet filters at their local site to

ensure enhanced security requirements. Figure 1 depicts the current network infrastructure. The remaining sites CINECA, CSC and EPCC which currently use the old 1 Gb/s phase 1 infrastructure because of infrastructural and technical delays will be connected within the next three month. This will complete the phase 2 DEISA network installation, because it is currently not planned to connect ECMWF to the 10 Gb/s network directly.

Backup scenarios have been considered, but not yet installed, via cross boarder fibres (CBF) between neighbouring NRENs. A backup scenario using GÉANT2 fibre footprint is under discussion also. A backup path being essential for most of the production Grid infrastructures today has not been of first priority in DEISA until now. The current design was chosen to minimize network costs, but should allow maximum outcome. A potential fiber cut or interface failure was accepted, because this implied a connectivity loss of one DEISA site only. This risk seemed to be acceptable since the availability of today's carrier networks like GÉANT2 is much higher than those of the connected HPC machines. This situation may change in the future. Additionally the partner supercomputer systems are working well as stand alone systems also, if a network link is not available. They simply don't schedule jobs which require remote data if the necessary links are down.

Nevertheless in future Grid environments computer systems will be tied together even more so that backup paths become necessary. Therefore the design has been planned to allow a redesign of the central DEISA network and install potentially three switches connected by a high performance ring infrastructure across Europe (30-40 Gb/s) and connect the DEISA sites to two of these by cross boarder fibres (dual homing). This allows easy rerouting in case of link disruption because of parallel connectivity. This design would also enhance overall network bandwidth between sites by load balancing traffic or splitting traffic dependent on destination. A schematic overview of a possible future backup design is depicted in figure 2.

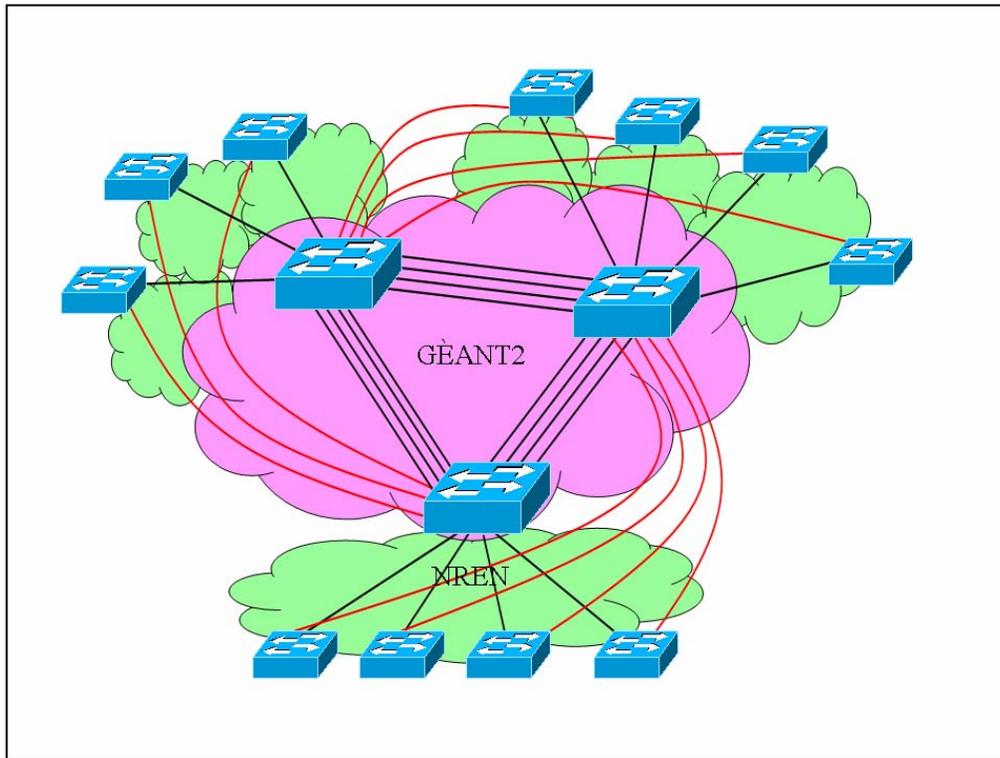


Figure 2: A schematic view of a possible future DEISA Backup Configuration

3 Managing a 10 Gb/s European network infrastructure

The management of a dedicated private 10 Gb/s WAN infrastructure connecting supercomputers across Europe does not differ much from general LAN management. The main difference comes from the fact that multiple local organisations, NRENs, national fibre providers and international network providers are involved. Therefore a multi domain network management system had to be developed, that allows to operate the infrastructure in a 24/7 fashion, to identify errors, to provide solutions and to operate an advisory service from the network point of view. Integrating the DEISA backbone into the normal network management procedures is a permanent activity. Though the infrastructure can be seen as only another small sub network of the local network that has to be monitored, the WAN character of this environment requires additional activities. In a local environment a link interruption can be handled by controlling cables and switch interfaces, changing equipment and retesting functionality. In principle this is the same for an European WAN. The difference arises when the error is located somewhere within the infrastructure. A fibre cable may have been cut on the way from Barcelona to Frankfurt. Because of the “simple” nature of the DEISA network to be managed, consisting of one central switch and only edges at the DEISA partner sites, the intermediate equipment will be hidden. There are many fibre sections in between where the problem could be located. Patch panels at NREN sites, GÉANT2 and local sites, WAN DWDM equipment including intermediate wavelength amplifiers may be the source of the problem. This implies involvement of many staff members of different organizations (local site, cable provider, lower layer service providers, NRENs, GÉANT2, DEISA network management staff) acting together for problem solution. This personal

communication infrastructure had to be organized, centralized, managed and updated continuously.

Beneath normal network monitoring based on standards like SNMP, RMON, Netflow/ipfix and syslog providing information about network utilization, logging, accounting, alarming, error recovery and problem determination, active probing procedures based on IPERF [6] and ping have been installed to log available TCP and UDP throughput dependent on system and network load in a 24 hours 7 days a week fashion. A special web site has been set up which provides this information to DEISA administrators and users. This network management page can be accessed at

<http://wwwnet.deisa.fz-juelich.de>

from authorized hosts and subnets only. As argued above a close collaboration between the DEISA network team and staff members of the NRENs and GEANT2 as well as a close interaction with the administrators of the supercomputer systems guarantees optimum performance of the DEISA network to meet the needs of the user communities. A close interaction with the DEISA "Operation team", responsible for any software or hardware changes within DEISA as well as any issues related to an undisturbed operation of the supercomputer systems, assures the production quality of the overall DEISA environment.

4 Operation experiences, Throughput values, et al.

The installation, configuration and operation of a European WAN will not happen without any pitfalls, surprises and new experiences. The lessons learned, local network and system configuration parameters used, throughput values seen between benchmark applications from the network point of view (Iperf) and real user applications (GPFS, GridFTP, ...) [7][8] as well as user acceptance and usability will be of overall importance for other organisations planning to install similar infrastructures. It has not been expected to measure a 10 Gb/s throughput for a single stream application. But the question arose what the overall aggregated performance of the underlying network will be and how a single flow/application can be optimized.

Optimized network throughput can be achieved by applications only if the systems have been configured accordingly. A good source for configuration information can be found at the "advanced networking" web page of the Pittsburgh Supercomputing Center [9]. Mathis et al. provide detailed information which networking parameters to alter and adjust. Unfortunately these settings change from OS version to OS version. Though most of them are obvious to be altered, others depend on settings at communication partner's site. They will have an effect only if the corresponding site changes parameters also. Others depend on remote setting in the way that they increase performance if also set on partners site and degrade performance if not. Since many of these parameters can be set only system wide and not interface specific they influence also local communications. E.g. high default buffer settings allow a large bandwidth delay product, but are inefficient if thousands of local connections have to be initiated. The values proposed within DEISA environment are shown in figure 3.

| | | |
|-----------------|-----------------------|----------------------|
| ipqmaxlen=2048 | tcp_ecn=0 | udp_pmtu_discover=1 |
| rfc1323=1 | tcp_mssdfmt=1460 | udp_recvspace=655360 |
| rfc2414=1 | tcp_newreno=1 | udp_sendspace=655360 |
| sack=1 | tcp_nodelayack=0 | use_isno=0 |
| sb_max=20971520 | tcp_pmtu_discover=1 | |
| thewall=1572864 | tcp_recvspace=2621440 | |
| | tcp_sendspace=2621440 | |

Figure 3: Optimised network options used at the DEISA supercomputer sites

After having installed and tuned the DEISA 10 Gb/s network we measured performance of the underlying infrastructure. It has been obvious that we will not see throughput values measured in an IBM laboratory test environment which has been about 5.6 Gb/s with an 10 Gb/s interface (10GBaseLR, Typ 7040-5719) in an P5-system with RIO3 drawers. The performance results have shown that throughput varies dependent on operating system version and installed hardware, local configurations and load of the system (we are running in production). So we have seen at FZJ about 4.8 Gb/s IPERF throughput to *localhost* which implies only using internal CPU and no network hardware, 4.3 Gb/s between 2 nodes of the FZJ system using the supercomputer internal IBM Federation switch and about 3 Gb/s between 2 nodes of the FZJ system using the 10 Gb/s GE interfaces connected by a CISCO Catalyst 6500 switch. Similar results can be measured at RZG, but 20-25 % less, which arises from slower CPUs within the RZG systems. These values can also be verified by testing IPERF throughput between a node in FZJ and another in RZG, both connected by a 10 Gb/s Ethernet interface to the DEISA-network. Throughput values reach up to 2.4 Gb/s. Here we started the client at FZJ and the server at RZG with the following commands:

```
Client: root@fzj::>iperf -i1 -w16000K -f m -c rzg
```

```
Server: root@rzg::>iperf -w16000k -s -D
```

System buffers have been aligned accordingly and jumbo frames have been used. The bandwidth delay product requires about 10 Gb/s * 13 ms \approx 16 MB. As mentioned before the systems used for these tests have been in production, so the measured throughput values varied dependent on system load. Sometimes intermediate values of about 3.2 Gb/s could be seen checking throughput in 1 second intervals using the “-i1” option of IPERF.

Of course those performance values may not be seen in “real” applications. The standard applications mostly used within DEISA are GPFS and GridFTP. Both of them allow use of parallel streams to optimize throughput. Using 4 parallel sessions with GridFTP a file can be transferred with a throughput of about 2.7 Gb/s from local FZJ site’s /etc/dev0 to remote RZG site’s /dev/null. Using /etc/dev0 and /etc/dev/null bypasses any delays because of disc access.

5 Network design and its influence on applications

Several heterogeneous supercomputer systems with different operating systems and system internal network configurations and various LAN designs are used within the DEISA backbone infrastructure. The local connectivity to the DEISA edge switches differs also. I/O nodes are connected with one 1 Gb/s or multiple 1 Gb/s (channelling) or

one 10 Gb/s connection to DEISA. Dependent on these local network designs different throughput values will be measured, because the involved nodes are providing different services and are therefore strained with potential additional work load. The following figures describe some of these scenarios to be found within DEISA:

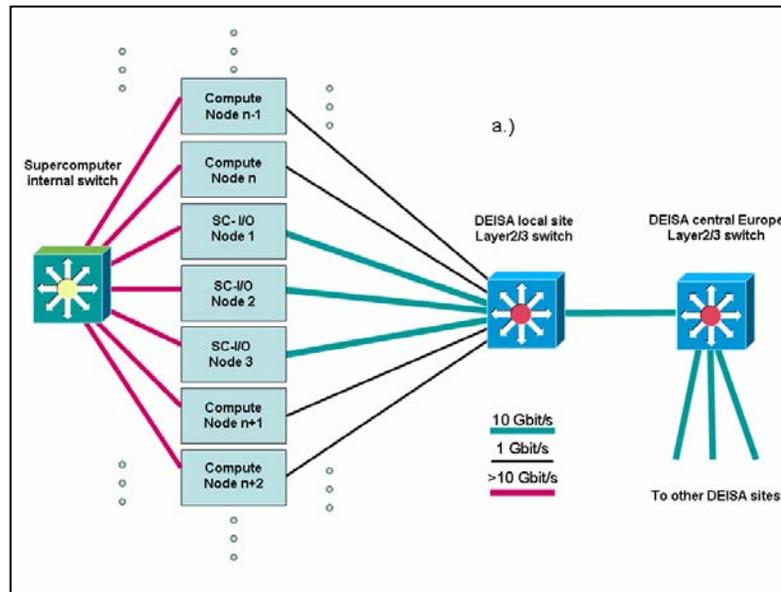


Figure 4 a

- Figure 4 a. describes a setup with compute nodes connected by 1 Gb/s and I/O nodes connected by 10 Gb/s for optimal access to data. The internal supercomputer switch is used for internal communications only. The setup allows many connections to be started in parallel with 1 Gb/s each. Whether a remote application can access the data stored at the I/O nodes with more than 1 Gb/s depends on the remote supercomputer network configuration.
- Within Figure 4 b. all nodes of the supercomputer system are connected with 1 Gb/s to the 10 Gb/s network infrastructure. The DEISA backbone allows many 1 Gb/s streams to cross the network in parallel. No stream can be initiated which exceeds the local 1 Gb/s link bottleneck. The internal network will not be used for DEISA purposes.

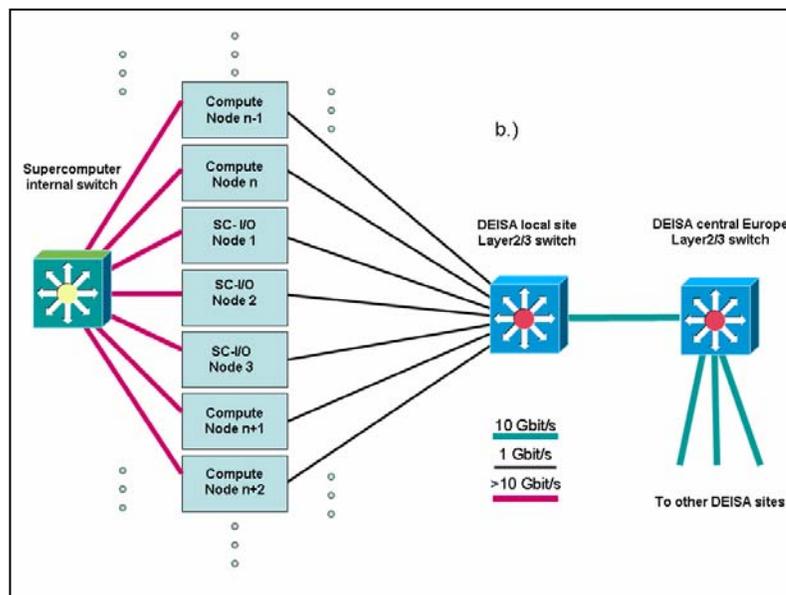


Figure 4 b

- Figure 4 c. describes a scenario where a special node provides a 10 Gb/s link from the supercomputer to the DEISA backbone working as a gateway node. All communications from I/O nodes and compute nodes must traverse the gateway node via the internal supercomputer network.

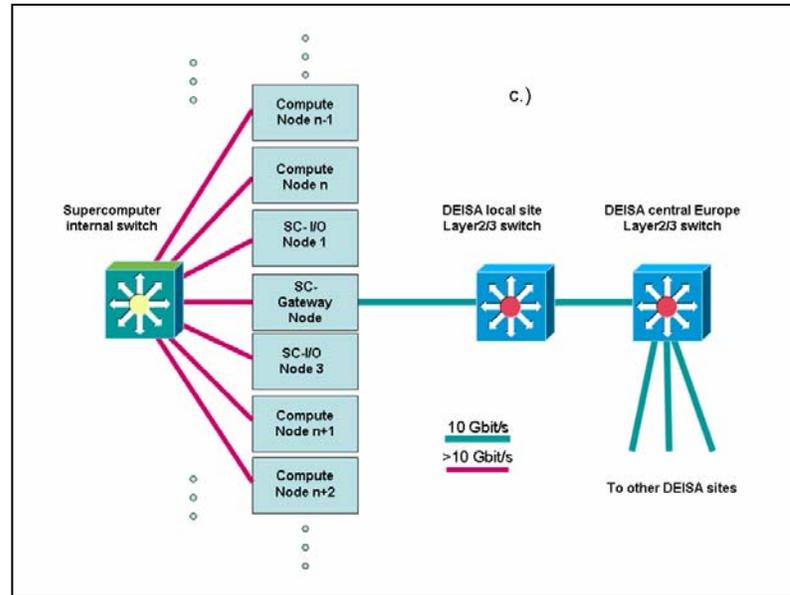


Figure 4 c

This setup allows an easy integration of the supercomputer system into the DEISA infrastructure and potentially enables a single 10 Gb/s application stream from every node to the DEISA backbone, but is highly dependent on the power of the gateway system, acting as a “software” router.

At FZJ we saw a P4 node consuming approximately one CPU per 1 Gb/s forwarding from internal Federation network to 10GE. That rate was not linear scaling to higher throughput and didn't reach the measured IPERF bandwidth numbers for 10GE usage without forwarding.

- The scenario in Figure 4 d. differs from 4 c. in the way that there is no special gateway node. A FrontEnd or Login node provides access to the supercomputer system. The same node will also be used for none DEISA purposes. Being the cheapest and simplest connectivity setup it implies disadvantages because of the heavy load which will be generated at the login node. Often this node will be also used for interactive purposes, so that users will not be really satisfied because of overload of the login node. Depending on the configuration this setup can be used to only access the system (really FrontEnd) which implies the supercomputer compute nodes are not directly reachable from DEISA or the

FrontEnd node also provides gateway functionality giving access to I/O and/or compute nodes.

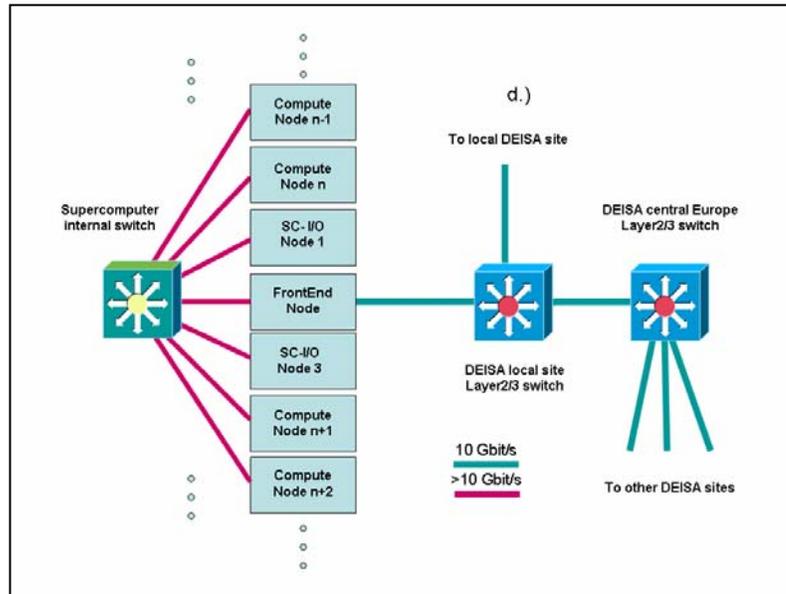


Figure 4 d

A more detailed technical view about an exemplary lambda setup at FZJ is shown in figure 5.). Here 2 I/O nodes are connected to the FZJ local DEISA edge switch with 10 Gb/s. A third node is configured as Gateway node and connected with 10 Gb/s also to DEISA. All other supercomputer nodes, named compute nodes are connected by 1 Gb/s Ethernet interfaces. For test purposes these nodes can be configured to directly communicate to DEISA or alternatively use the 10 Gb/s gateway node. Any mixed setup is also possible only dependent on local compute node configuration and FZJ local DEISA leaf switch setup. The DEISA edge switch has one 10 Gb/s Long Range Interface which is connected to the local NREN DFN DWDM equipment. Using DFN's national fibre footprint one special, meanwhile optical protected, wavelength is reserved for DEISA. This wavelength is terminated at the DEISA central switch in Frankfurt, Germany. All other sites have a similar setup. The only difference for non German sites is that their wavelengths are transparently prolonged through the European GÉANT2 network which has been sketched in the right lower corner of Figure 5. Figure 5 doesn't show the normal links to the local FZJ network where all the other FZJ hosts are located nor the "normal" Internet connection of FZJ, which can be access via the login node at the left lower corner of this picture.

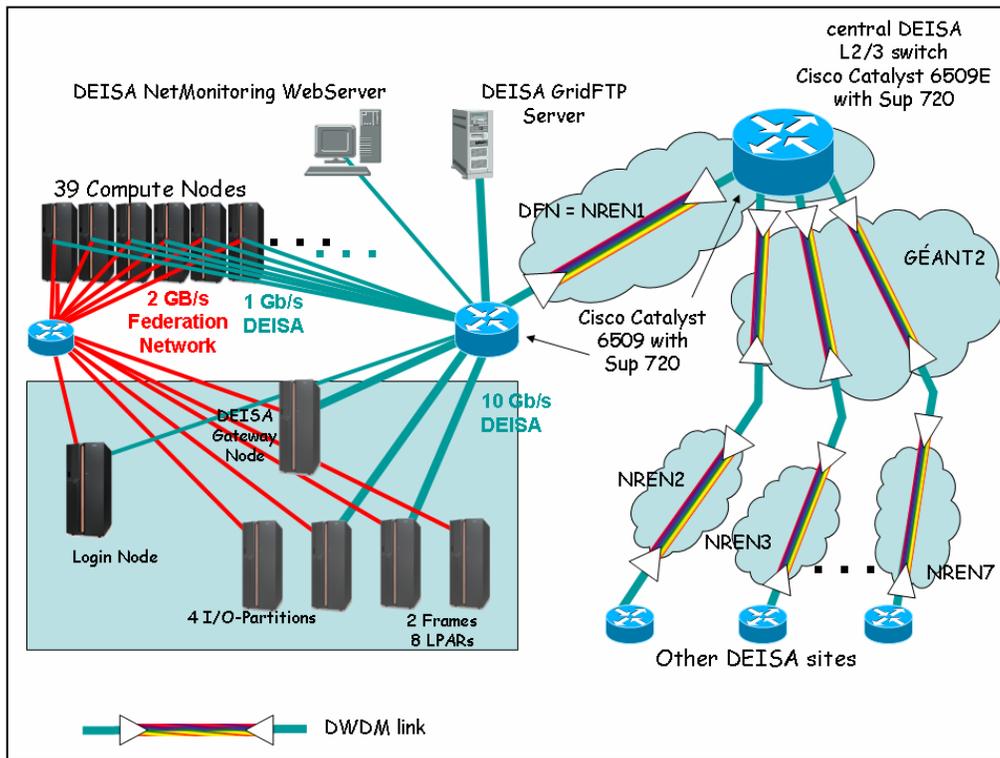


Figure 5: Schematic view of DEISA lambda setup

6 Lessons learned

The operation of an infrastructure like DEISA leads to new management problems not seen before. Managing a supercomputer system or a number of locally installed cluster systems differs heavily from an European supercomputer infrastructure where staff members dealing with the same problem are thousands of miles away. There is no short cut, going to the office next door, just checking if we agree on some option settings within a software component. Within a virtual organization every small modification has to be checked by all partners over and over again. Installing new software components requires checking with all participants, if any dependencies exist. Scheduling of tasks, installations, maintenance, network infrastructure changes and others have to be agreed on. Often a task needs much more of time than estimated. Someone has to deal with those issues. DEISA has identified those issues and has introduced an "Operations Team" which is working as an interface between the different service activities and handling all these interdependencies. Since creation of this new team, the network team is in close contact with this group which has made interaction with the other service activities much easier.

Though all these things can be handled by e-mail mostly, it is nevertheless mandatory to have regular phone and video conferences, writing minutes and checking for completion of tasks. Additionally it is often necessary to have agreed on strict rules for processing if any disagreements arise. Those dissents are mainly found among others in security policy issues, scheduling of software installation and upgrades, budget issues for needed components and other issues. From the network point of view the optimal network setup requires extensive discussions. Often simple network redesigns will lead

to time-consuming discussions. Local priorities often contradict optimal solutions. Especially in a multi to multi network scenario (DEISA has 11 sites) a design has to be chosen, which arises from the lowest common denominator. Things have to be configured that allow a good performance for all sites, which may not be the optimal solution for every connection setup.

7 Future enhancements and co-operations

The eSA1 activities could not be started before November 2006, because of the missing eDEISA budget. Therefore contracts could not be signed and equipment could not be ordered in advance.

The main problems that exist in connecting the remaining sites are local provider solutions and delays because of equipment deliveries. Especially GÉANT2 needs about 12 weeks to order and provide adequate DWDM equipment, to test the new fibers and wavelengths and to configure NREN-GÉANT-NREN links through Europe. It is assumed that the remaining DEISA partners can be connected to the 10 Gb/s backbone within the next 3 months.

DEISA is aware of all the activities around Europe and measures itself only upon its mission to enable new science in Europe. Many new activities will come up with the European FP7 program. DEISA has to find its position in these upcoming activities. DEISA is open for new grand challenges and is prepared to be part of the future European HPC infrastructure. Currently T1 High Performance Centres (HPC) are integrated into DEISA, but the upcoming European T0-HPC centres will be of major importance to DEISA. DEISA sees its role as an integrating part of those centres being the global player. Even if T0 centres should not be fully integrated into the DEISA infrastructure, a close collaboration will be indispensable. Therefore DEISA plans to integrate these sites into its backbone infrastructure implying further evolution of the network in future.

DEISA feels confident that its challenging goals are highly dependent on the future development of future national and international networks. Future HPC centres may be able to operate and serve its users only if an adequate networking connectivity can be provided. Researchers are spread across Europe, but distances become shorter. These scientists work together more and more looking for resources where ever possible. In future it will become irrelevant where the supercomputer system is located. Connectivity will be the driving force.

Cost-efficient high-performance interconnects need to be provided. Future bandwidth-on-demand services will allow the exact reservation of network links with appropriate bandwidth, quality and availability. GÉANT2 and the European NRENs have done a good job until now, but they will have to evolve further on. A European world-class, high-bandwidth, multi protocol network infrastructure will push forward science to new dimensions.

Three years of successful DEISA operation have shown that the concept implemented in DEISA proceeded very well. This does not preclude that organizational structures of DEISA may change over time. But the general idea of DEISA will sustain. The next steps within DEISA will be to establish an efficient organization embracing all relevant HPC organizations in Europe. Being a central player within European HPC, DEISA intends to contribute to a global e-Infrastructure for science and technology furthermore. Integrating

leading supercomputing platforms with Grid technologies and reinforcing capability with shared petascale systems is needed to open the way to new research dimensions. The new DEISA 10 Gb/s European network infrastructure provides the fundament for the DEISA supercomputer infrastructure through which DEISA will pave the way for further scientific computing. The DEISA supercomputing infrastructure, inviting leading scientific research areas to operate on future grand challenges, provides a vision of future European Grid computing.

8 Summary

eSA1 is on its way to a full 10 Gbit/s star like network infrastructure. Due to the delays mentioned above we are behind expected time schedules, but these problems have been solved within the last month. Tasks to be done in the near future are the transitions of the remaining partners from 1 Gbit/s to a 10 Gbit/s connectivity as well as the upgrades needed for the integration of the monitoring of the new infrastructure. These tasks will be started soon.

9 References and Applicable Documents

- [1] Distributed European Infrastructure for Supercomputer Applications, <http://www.deisa.org>
- [2] GÉANT home page, <http://www.geant.net>
- [3] GÉANT2 home page, <http://www.geant2.net/>
- [4] GÉANT - GÉANT/Dante description of the Premium IP service, <http://www.dante.net/server/show/nav.00700a003>
- [5] Y. Rekhter, B. Moskowitz, D. Karrenberg, G.J.de Groot, E. Lear, Address Allocation for Private Internets, <ftp://ftp.rfc-editor.org/in-notes/rfc1918.txt>
- [6] lperf - Version 2.0.1, The National Laboratory for Applied Network Research (NLANR), Distributed Application support team, <http://dast.nlanr.net/Projects/lperf/>
- [7] GPFS: A Shared-Disk File System for Large Computing Clusters, F.Schmuck, R.Haskin, Proceedings of the Conference on File and Storage Technologies, 28–30 January 2002, Monterey, CA, pp. 231–244., http://www.almaden.ibm.com/StorageSystems/file_systems/GPFS/Fast02.pdf
- [8] W. Allcock, GridFTP: Protocol Extensions to FTP for the Grid, Open Grid Forum document, <http://www.ogf.org/documents/GFD.20.pdf> and I. Mandrichenko, GridFTP Protocol Improvements, Open Grid Forum document, <http://www.ogf.org/documents/GFD.21.pdf>
- [9] M.Mathis, R.Reddy, J.Mahdavi, Advanced Network computing – Enabling High Performance Data Transfers, <http://www.psc.edu/networking/projects/tcptune/>, Nov 2006

10 List of Acronyms and Abbreviations

| | |
|-------------------------|---|
| BPDU | Bridge Protocol Data Unit - A special message type within the Spanning Tree Protocol for switches |
| BSC | Barcelona Supercomputing Center, Barcelona, Spain |
| CINECA | Consorzio Interuniversitario, Bologna, Italy |
| CBF | Cross Boarder Fibres – direct fibre links between neighbouring countries not provided and managed by DANTE |
| CSC | Finnish Information Technology Centre for Science, Espoo, Finland |
| DEISA site | partner site of the DEISA consortium |
| DWDM | Dense Wavelength Division Multiplexing |
| ECMWF | European Centre for Medium-Range Weather Forecasts , Reading, UK |
| EPCC | Edinburgh Parallel Computing Centre, Edinburgh, UK |
| FZJ | Forschungszentrum Jülich, Jülich, Germany |
| GÉANT | multi-gigabit pan-European data communications network administrated and operated by Dante |
| GÉANT2 | the seventh generation of pan-European research and education network, successor of GÉANT |
| GÉANT2 footprint | The fibres leased or owned by DANTE which are building the GÉANT2 network |
| HLRS | High Performance Computing Center Stuttgart, Stuttgart, Germany |
| IDRIS | Institut du Développement et des Ressources en Informatique Scientifique, Orsay, France |
| JumboFrames | Ethernet frames concerning standard IEEE 802.3 are 1518 Byte long. Jumbo frames are not standardized Ethernet frames larger than that (mostly 9kB) |
| LRZ | Leibniz Computing Centre of the Bavarian Academy of Sciences and Humanities, Garching, Germany |
| NREN | National Research and Education Network, e.g. DFN in Germany, RENATER in France, SURFnet in The Netherlands |
| Local provider | This may be the NREN where a DEISA site is connected to directly (the DEISA site is a NREN PoP) or a fiber company which provides fibers to the first NREN PoP of the local DEISA site. |
| PREMIUM IP | Service provided by DANTE with special service level agreements (SLA) |
| PoP | Point of Presence of a NREN. That is the location where a NREN has its equipment installed and to which a site has to connect to, to get access to the world wide internet. |
| RZG | Garching of the Max Planck Society, Garching, Germany |
| SARA | SARA Computing and Networking Services, Amsterdam, The Netherlands |