

Analysis of internal network requirements for the distributed Nordic Tier-1

G Behrmann¹, L Fischer², M Gamst³, M Grønager¹ and J Kleist⁴

¹ NDGF - Nordic DataGrid Facility, Kastruplundgade 22(1), DK-2770 Kastrup

¹ NORDUnet, Kastruplundgade 22(1), DK-2770 Kastrup

³ Technical University of Denmark, Department of Management Engineering,
Produktionstorvet, DTU - Building 424, DK 2800 Kgs. Lyngby

⁴ NDGF and Aalborg University, Department of Computer Science, Selma Lagerlöfsvej 300,
DK 9220 Aalborg SØ

E-mail: behrmann@ndgf.org, lars@nordu.net, gamst@man.dtu.dk, gronager@ndgf.org,
kleist@ndgf.org

Abstract. The Tier-1 facility operated by the Nordic DataGrid Facility (NDGF) differs significantly from other Tier-1s in several aspects: It is not located at one or a few locations but instead distributed throughout the Nordic, it is not under the governance of a single organisation but is instead build from resources under the control of a number of different national organisations. Being physically distributed makes the design and implementation of the networking infrastructure a challenge. NDGF has its own internal OPN connecting the sites participating in the distributed Tier-1. To assess the suitability of the network design and the capacity of the links, we present a model of the internal bandwidth needs for the NDGF Tier-1 and its associated Tier-2 sites. The model takes the different type of workloads into account and can handle different kinds of data management strategies. It has already been used to dimension the internal network structure of NDGF. We also compare the model with real life data measurements.

1. Introduction

Dimensioning the network for a Tier-1 is always a challenge, but in the case of the Nordic Tier-1 operated by NDGF, there is the extra challenge that the Tier-1 is distributed. The NDGF Tier-1 consists of the 7 biggest Nordic compute centers, dTier-1s, with associated Tier-2 resources as far away as Slovenia. Resources (storage and computing) are widely scattered with a few central services. This give a lot of advantages in redundancy, especially for 24x7 data taking as reported in [1]. Figure 1 shows the centers participating in the Nordic countries including the resources they have available as of Q2 2009.

NDGF has uses an internal OPN between all dTier-1 sites and the Slovenian Tier-2, other Tier-2 sites are connected via the national research network. Figure 2 shows the how NDGF sites are interconnected with special emphasis on Sweden; red lines depict dedicated private network lines and magenta lines depict public lines. As it can be seen from the figure, the main network infrastructure forms a star. This give us the advantage that it is easy to calculate the load on the links from the central NDGF router to each country; one just need to consider the resources in a country as a single site.

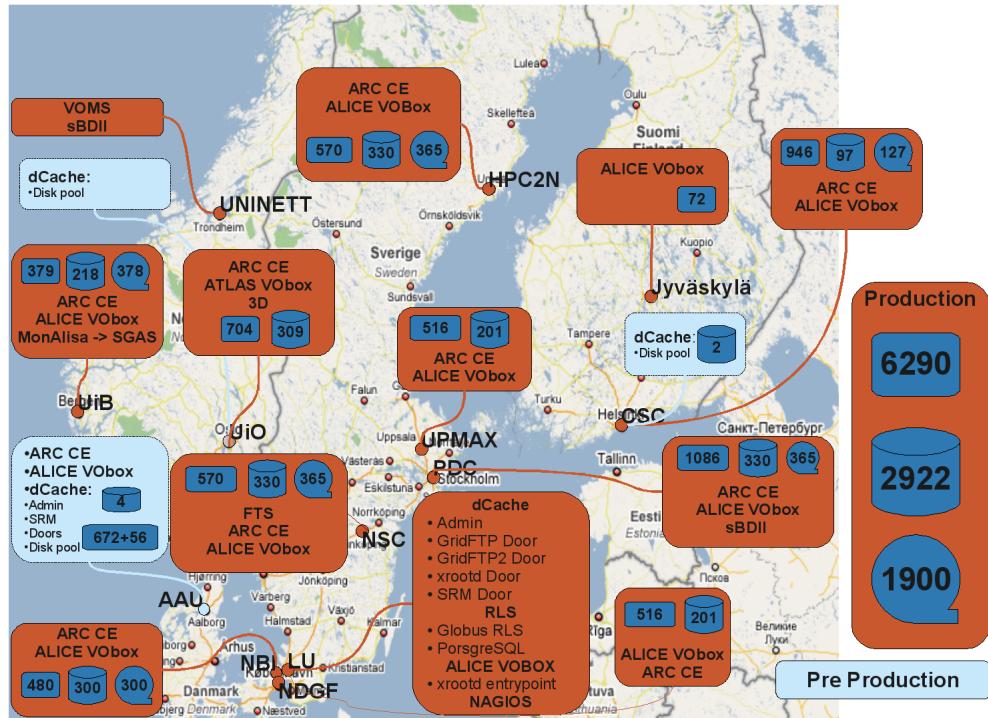


Figure 1. NDGF Distributed Storage and Computational setup.

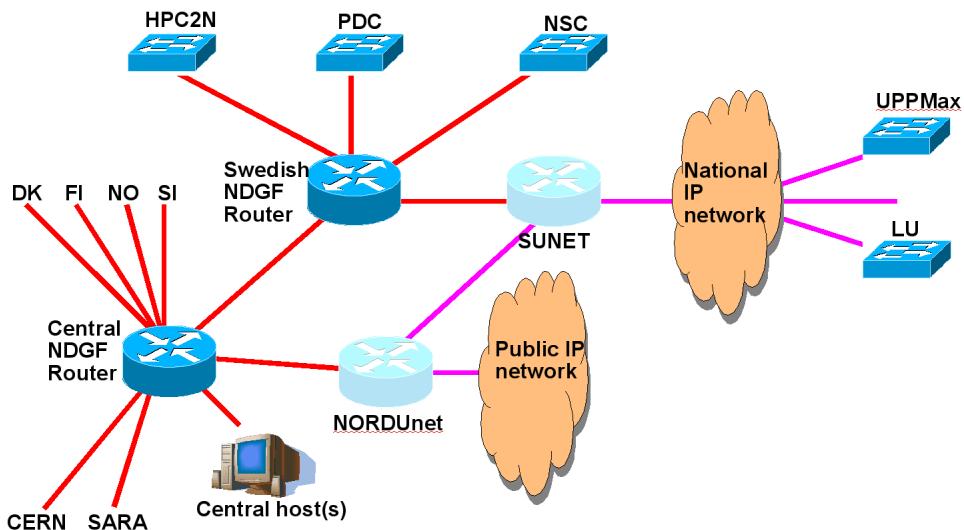


Figure 2. NDGF Network Layout.

2. Calculating network need

The aim of with the model is to get a estimate of what the bandwidth requirements is. For links within the Tier-1s we are interested in knowing if the 10 Gbps dedicated lines are sufficient and

for the public links to the Tier-2s we are first of all interested in knowing if the load can be carried without disturbing other uses of the network. We will try and construct a model, where the network load is driven by the amount of available compute power, assuming that there is enough storage available in the system.

2.1. Basic assumptions

Below we present that basic assumptions in the model.

- All worker nodes are occupied up to their efficiency, i.e. we assume that there is enough jobs and that all data is available in the system.
- Data is randomly and uniformly distributed over all storage sites, i.e. the chance that a piece of data is available at a specific storage site is proportional to the size of the storage site in relation to the total amount of storage in the system.
- The characteristics of different job types is known, i.e. we know in advance how much data a job consumes and generates and we know how many CPU resources are required for each job type.
- The job mix at a site is known in advance, i.e. we know how many percent a specific job type spends of the available compute time.
- Jobs are spread temporally uniformly a site. This means that we don't have bursts of a specific job type.
- Traffic flows directly between the compute site where a job is executed and the storage site. No intermediate servers are involved. This is in fact the case for the NDGF setup.
- The caching mechanism in the ARC grid middleware [2] is not taken into account. The ARC middleware employed by NDGF for its ATLAS computations includes a caching mechanism, that can quite significantly reduce network traffic. Modelling the impact of the caching mechanism is rather difficult without any empirical evidence on what effect the cache has on different job types.

Later in this paper we will discuss what impact changes to those assumptions will have on the model.

2.2. Site characteristics

For a site s the we assume to know the following characteristics:

- Amount of tape installed: T^s .
- Amount of disk installed: D^s .
- Tier- i compute resources: C^{s_i} .
- Tier- i efficiency: $e^{s_i} \in [0; 1]$.

Some NDGF sites act as both Tier-1 and 2 centers, this is why we allow a site to have a number of compute resources and efficiencies.

We let T and D denote the total amount of tape and disk in the system respectively.

2.3. Job characteristics

Just as for sites we need to know the characteristics of the jobs that are executed. NDGF runs a number of different job types, for each job j we assume to know:

- Amount of CPU seconds to run a job: R_j .
- Amount of data read from disk while executing: DI_j .
- Amount of data read from tape while executing: TI_j .

- Amount of data written to disk while executing: DO_j .
- Amount of data written to tape while executing: TO_j .

Each site runs a special jobmix - i.e. a job type j is supposed to occupy a certain fraction of the available compute time. We let J denote the total set of job types in the system. Let $f_j^{s_i}$ denote the fraction for job type j on the Tier- i resources at site s (we assume that $\sum_{j \in J} f_j^{s_i} = 1$).

We can now calculate the amount of data that needs to be read from and written to disk and tape at a site s to keep its worker nodes occupied.

Each job j needs DI_j data from disk and runs for R_j CPU seconds (this could be any general measurement of CPU performance like KSI2K or HEP2006). Therefore DI_j/R_j denotes the amount of data a job j requires per CPU second. We then just need to multiply with the amount of CPU resources available to that job type; for a Tier- i resource this number is given by $f_j^{s_i} e^{s_i} C^{s_i}$, i.e. the fraction of the resource that runs jobs of type j times the efficiency times the total amount of computational resources.

$$DI_C^s = \sum_{j \in J} \sum_{i \in \{1..2\}} f_j^{s_i} e^{s_i} C^{s_i} \frac{DI_j}{R_j}$$

Similarly we can calculate values for tape read TI^s , disk write DO^s and tape write TO^s :

$$\begin{aligned} TI_C^s &= \sum_{j \in J} \sum_{i \in \{1..2\}} f_j^{s_i} e^{s_i} C^{s_i} \frac{TI_j}{R_j} \\ DO_C^s &= \sum_{j \in J} \sum_{i \in \{1..2\}} f_j^{s_i} e^{s_i} C^{s_i} \frac{DO_j}{R_j} \\ TO_C^s &= \sum_{j \in J} \sum_{i \in \{1..2\}} f_j^{s_i} e^{s_i} C^{s_i} \frac{TO_j}{R_j} \end{aligned}$$

2.4. Bandwidth requirements

At a site s , parts of the data can be read and written locally. If we assume a uniform distribution of data over the sites, the part that can be read and is written locally corresponds to the fraction of storage available at s in relation to the total amount of storage in the systems.

$$BI_C^s = DI_C^s \frac{D - D^s}{D} + TI_C^s \frac{T - T^s}{T}$$

and similar for output.

Furthermore, other sites will read and write to the disk and tape systems at a site s . Again the amount corresponds to the relation between the installed disk and tape capacity at s and the total installed capacity.

$$BI_O^s = \sum_{t \in S \setminus \{s\}} \left(DO_C^t \frac{D^s}{D} + TO_C^t \frac{T^s}{T} \right)$$

Finally, we need to take traffic external to NDGF into account. Again the same arguments as before applies, and the traffic to a site is the fraction of the total traffic BI_E that corresponds fraction of the total storage that a sites has.

$$BI_E^s = BI_E \frac{D^s}{D}$$

And the input bandwidth requirements for a site s becomes $BI^s = BI_C^s + BI_O^s + BI_E^s$. For bandwidth out of a site, we derive a similar formula

$$BO^s = DO_C^s \frac{D - D^s}{D} + TO_C^s \frac{T - T^s}{T} + \sum_{t \in S \setminus \{s\}} \left(DI_C^t \frac{D^s}{D} + TI_C^t \frac{T^s}{T} \right) + BO_E \frac{D^s}{D}$$

The bandwidth requirement for a site is then the maximum of BI^s and BO^s .

This model will as output give us the average network throughput at a site if all compute resources are occupied up to their efficiency with a random mix of jobs. It will not take burst into account, neither will it include any overhead caused by transport protocols.

3. Main Results

Table 1 shows the outcome of using the model on the current NDGF setup. In Appendix A we list the values for the job types as well as the fractions expected each job type should take up on a Tier-1/2 site respectively. As it can be seen from Table 1 are most of the results within what we can expect to carry over a 10 Gbps line, the only worrying link is the link between the NDGF central router and Sweden. For the Swedish link a more carefull analysis based on real life measurement will need to be carried out.

Site/Country	Network load
DCSC/KU	0.8 Gbps
Denmark	0.8 Gbps
CSC	0.5 Gbps
Jyv	0.0 Gbps
Finland	0.5 Gbps
UiB	0.8 Gbps
UiO	0.7 Gbps
Norway	1.3 Gbps
HPC2N	1.9 Gbps
LUNARC	0.3 Gbps
PDC	1.2 Gbps
NSC	1.0 Gbps
UPPMAX	0.3 Gbps
Sweden	3.8 Gbps
PIKOLIT	0.6 Gbps
Slovenia	0.6 Gbps

Table 1. Network load between countries central routers and the NDGF main router.

4. Discussion

The results presented here can only be considered as a first approximation at what kind of network load NDGF can expect internally. Some of the assumptions behind the model can rightfully be criticised for being too simple. Especially the assumption on the uniform distribution of job types over time, is questionable. One way to deal with that assumption would be to only consider the job type that causes the highest network load. This would make the model a better fit for worst case loads. However, it will be more important to take the caching mechanism of ARC into account, as this mechanism has been reported to have a significant impact on how many times popular files are downloaded to a site. In order to extend the model to take caching into account, a more in-depth analysis of the caching mechanism of ARC need to be performed first.

At the time of writing it is not yet possible to compare the complete model to real world measurements. However, a few observations can be made: i) the limiting factor is a the moment not lack of network bandwidth, it is in cases observed so far caused by lack of bandwidth between the compute element and the network. This issue is solvable by upgrading the hardware at sites. ii) the assumption about a uniform mix of jobs does not hold. We observe that jobs of certain types comes in bursts. iii) The ARC caching mechanism has a dramatic effect on the amount of data transferred. This has been observed by correlating the load on the dCache system with the amount of data needed to carry out computations. In cases where data are cached, the load on dCache decreases.

We are currently investigating the deployment of the monitoring services needed to perform better measurements of the used bandwidth and to have the necessary information to test the validity of our model.

References

- [1] Field L, Grønager M, Johansson D and Kleist J *Proceedings of CHEP'09* (Submitted for publication)
- [2] Ellert M, Grønager M, Konstantinov A, Kónya B, Lindemann J, Livenson I, Nielsen J, Niinimäki M, Smirnova O and Wänänen A 2007 *Future Generation Computer Systems* **23** 219–240

Appendix A. Job type descriptions

The table presented here forms the basis for the calculation of the percentage of time that is spent on each job type at a site. The numbers cannot be used directly as is, since not every site support the same VOs.

Job name	Tier-1	Tier-2	Run time	Disk in	Disk out	Tape in	Tape out
ALICE analysis	20%	50%	1	1000	10	0	0
ALICE recon	40%	0%	5	10	100	1000	0
ALICE MC	40%	50%	15	10	10000	0	0
ATLAS analysis	20%	50%	1	100	100	0	0
ATLAS recon	40%	0%	1	10	100	1000	0
ATLAS MC	40%	50%	12	100	500	0	0
CMS analysis	20%	50%	1	100	100	0	0
CMS recon	40%	0%	2	100	100	2000	0
CMS MC	40%	50%	12	100	500	0	0