

File ID 164787
Filename Summary

SOURCE (OR PART OF THE FOLLOWING SOURCE):

Type Dissertation
Title A semantic model for complex computer networks : the network description language
Author J.J. van der Ham
Faculty Faculty of Science
Year 2010
Pages viii, 154
ISBN 978-90-814160-2-3

FULL BIBLIOGRAPHIC DETAILS:

<http://dare.uva.nl/record/331711>

Copyright

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use.

Summary

Communication over computer networks is an important part of our society today: we make phone calls, send emails, and surf the web. All these processes are enabled by the physical infrastructure of wires and fibers in the ground, combined with networking devices that communicate electronically.

Scientific research has also adapted to make use of the increased network connectivity. Researchers are sharing data sets, computing infrastructure and specialized equipment, all through the network. However, more and more scientific applications need better quality of services than the regular packet-switched Internet can offer. Such applications may produce so much traffic that if they use the regular Internet, they fail to run smoothly or disrupt other Internet traffic. These applications require dedicated network connections.

In the last few years a number of national research and education networks have moved to hybrid network infrastructures based on optical networking. The term ‘hybrid network’ means that the network carries the regular Internet traffic, and a different part is used to provision lightpaths, dedicated connections, to support researchers.

Lightpaths can span multiple administrative domains. A typical network connection between two universities first crosses a campus network, then a national research network, then goes through an international peering to another research network, then another campus network and finally to an internal network within a building. Different persons and organisations administrate all these networks.

Ideally, the scientific applications that need to move large amounts of data, or require a smooth network service simply ask for a lightpath and it is provided automatically. Currently, the configuration of a lightpath is performed manually and can take up to several weeks. The path through the topology must be specified and details communicated to all the network providers before the

lightpath can be configured.

The first part of this thesis describes the Network Description Language, a model for describing complex network topologies and technologies. The model defines a clear terminology for describing network topologies which can be linked to other descriptions and other resources. In the same way topology descriptions can be linked to descriptions of other networks, creating a distributed view of the global topology. These linked topologies can then be used to facilitate the configuration of lightpaths and optical networks.

Ideally all the network providers maintain the topology information so that there is a complete overview of possible connectivity and network details. However, network operators tend to be protective of detailed topology information, because of scalability, security, or policy reasons. Fortunately, it is also possible to share an aggregated view of the topology. This means that only some details of the network are published.

There are different levels of abstraction that can be used for topologies: from collapsing a whole network domain to a single point, to providing detailed information about the edge nodes and their internal connectivity. In part two of this thesis I describe an emulation to find out what kind of effect this aggregation has on the accuracy of inter-domain pathfinding.

From the results we can see that publishing detailed information about edge nodes and their internal connectivity shows almost no difference with pathfinding on unaggregated topologies. Collapsing a domain to a single node is clearly the worst performing solution. From the results we can conclude that with higher levels of abstractions, the accuracy of inter-domain pathfinding degrades significantly.