GeS 2007
German e-Science

# The WISENT Grid Architecture: Coping with Firewalls and NAT

G. Scherp, W. Hasselbring, J. Ploski

WISENT
Wissensnetz Energiemeteorologie

## The WISENT Grid architecture

The figure on the right side depicts the Grid architecture of our research project WISENT (http://wisent.d-grid.de). Blue resources exist at our project partners, and the green resources are available in the German Grid (D-Grid) infrastructure. The cluster at OFFIS and the storage at DLR, except for the DIMS archive, are dedicated D-Grid resources.

Except for the cluster at OFFIS, internal computing resources of each project partner are managed by Condor middleware. The DIMS archive can be accessed only through proprietary services. Dedicated D-Grid resources at OFFIS and DLR run middleware required by D-Grid.

All Grid resources at our project partners are externally connected via Globus Toolkit 4 to support data transfers via GridFTP and the RFT service as well as to enable access to external computing resources via WS-GRAM, which is also capable of accessing Condor pools.

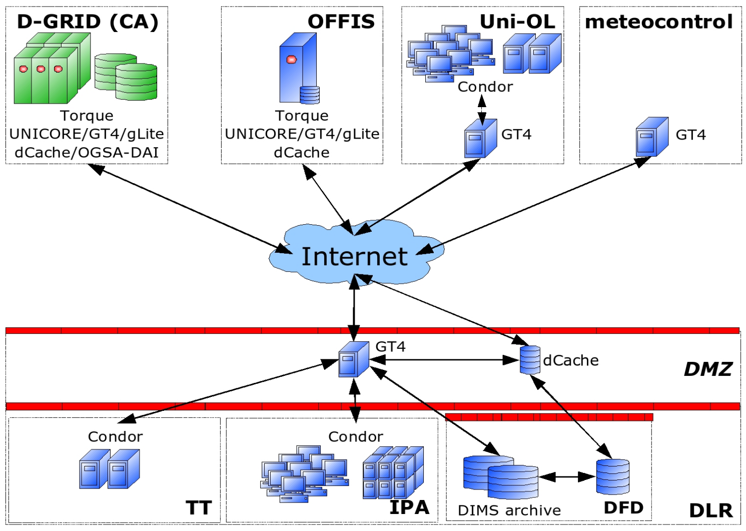

## Problems with firewalls and NAT

In the Grid architecture the communication between Grid software (Grid clients, Grid middleware, batch systems) is hampered by strong firewalls and security policies at DLR. The DLR network is divided into a demilitarized zone (DMZ) and an internal network, with the following communication characteristics:

- Most incoming connections into the internal network are completely blocked.
- Most outgoing connections are allowed without restrictions.
- Only incoming connections on static ports are normally allowed.

The resulting problems are:

- Components of Globus Toolkit 4 cannot simply use dynamic ports for data channels (GridFTP) and callbacks (WS-GRAM) [Wel06].
- Direct connections to a Condor pool or data storages in the internal network from Globus Toolkit 4 in DMZ are not possible.
- Callbacks of an external WS-GRAM service to clients in the internal network are blocked as unsolicited incoming connections.

The use of network address translation (NAT) causes similar problems as it usually blocks all incoming connections (due to technical reasons) and allows outgoing connections. Furthermore, the use of one external IP address for multiple internal hosts conflicts with the Grid security infrastructure (GSI), which demands that each host is reachable via a so-called fully qualified domain name (FQDN) [Wel06].

## A solution approach with "hole punching"

Recommendations for firewall configurations of the D-Grid integration project (DGI) [VG06] and existing workarounds between the DMZ and the internal network solve some of the above problems. However, general technical solutions for direct communication between Grid software in presence of firewalls and NAT are still missing.

A possible approach is to adopt the "hole punching" technology, which is already used in peer-to-peer (P2P) networks, when two clients are behind NAT systems and are unable to establish a direct connection because they lack globally unique IP addresses. Hole punching works with UDP (used by most P2P applications) as well as with TCP [For05]. In general, this technology relies on a broker (which has a globally unique IP address) that works as a mediator to establish a direct communication link between two clients. The following simplified example illustrates hole punching:

- Host A submits the request for communication with Host B to Host S
- Host S submits the public endpoint (port 4000) of Host A to Host B.
- Host B sends a packet via it's public-private endpoint (port 5000) to the public endpoint of Host A (4000). The NAT system of Host A rejects the packet, but a NAT session on B's side is now established ("a hole is punched"), ready to translate any packets from A's public endpoint (port 4000) to B's public endpoint (port 5000) into B's corresponding private endpoint.
- Host B notifies Host S and awaits an incoming connection from Host A.
- Host S submits the public endpoint of Host B (port 5000) to Host A.
- Host A uses its public-private endpoint (port 4000) to establish a connection with Host B's public endpoint (port 5000).

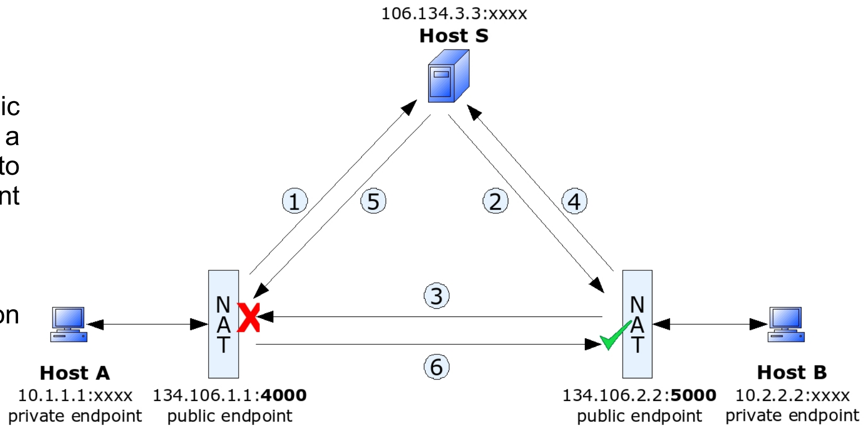

Beside NAT systems it is also possible to use hole punching to traverse firewalls. Thus, in the above figure NAT systems could be replaced with firewalls, separating hosts running Grid software. However, Grid software communicates via TCP and hole punching with TCP is more prone to errors due to specific firewall and NAT configurations [For05]. We implemented a prototype broker and clients to test feasibility of TCP hole punching in our context. The next step is to evaluate the concepts by extending components of Globus Toolkit 4.

[For05] *B. Ford. Peer-to-peer communication across network address translators, 2005. USENIX Annual Technical Conference.*
[VG06] *Gian Luca Volpato and Christian Grimm. Empfehlungen zur statischen Konfiguration von Firewalls im D-Grid. Technical report, 2006.*
[Wel06] *Von Welch. Globus Toolkit Firewall Requirements. Technical Report 9, Globus Alliance, 10 2006.*

OFFIS
OFFIS Institute for Information Technology
R&D Division Business Information Management (BI)

Carl von Ossietzky universität Oldenburg
University of Oldenburg
Software Engineering Group

SPONSORED BY THE
Federal Ministry of Education and Research