

Distributed OpenMI

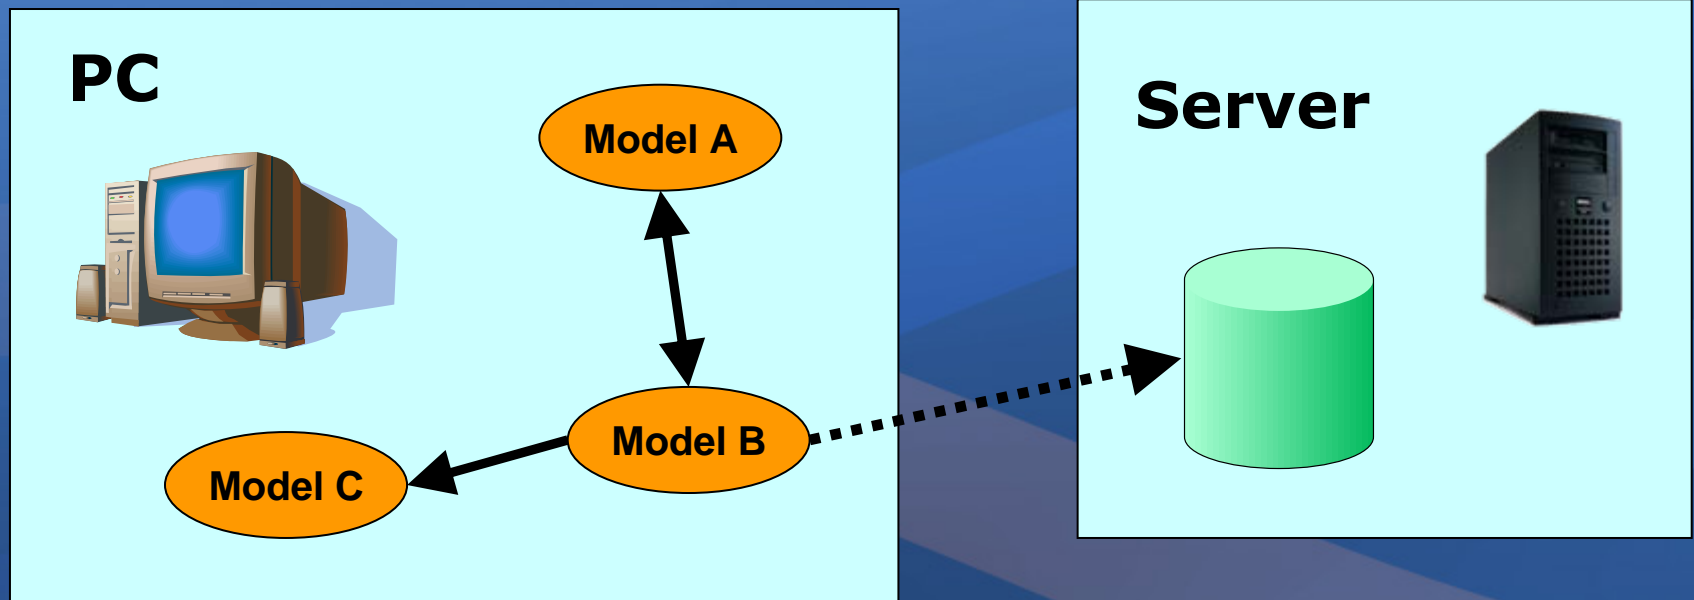
by Jan Čurn
(j.curn@dhi.cz)

DHI




WATER • ENVIRONMENT • HEALTH

- OpenMI models reside on a single computer, in same application
- Linkage and computation in single thread
- Linkage of models running on separate computers VERY difficult



- Allows linkage of models running on different computers
- Transparent to OpenMI-compliant models
- Simple to use

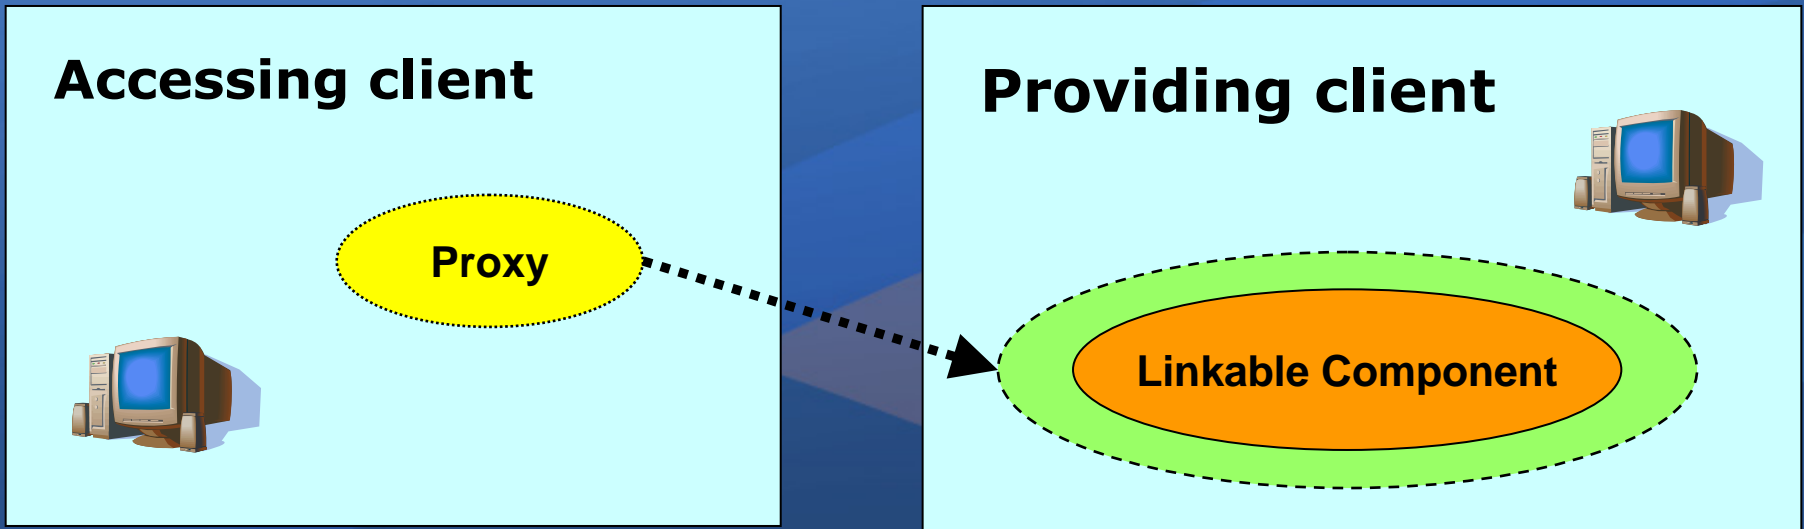


- Remote access to large time-series databases
- Deploy the simulation on the desktop, run it on a dedicated machine
- Efficient usage of simulation software licenses
- Link models from different providers without a need for moving data
- Linux vs. Windows, Java vs. .NET interoperability


How it works ?



- Accessing client vs. providing client
- Communication over network subsystem (using RPC protocol)
- YA-RPC protocol

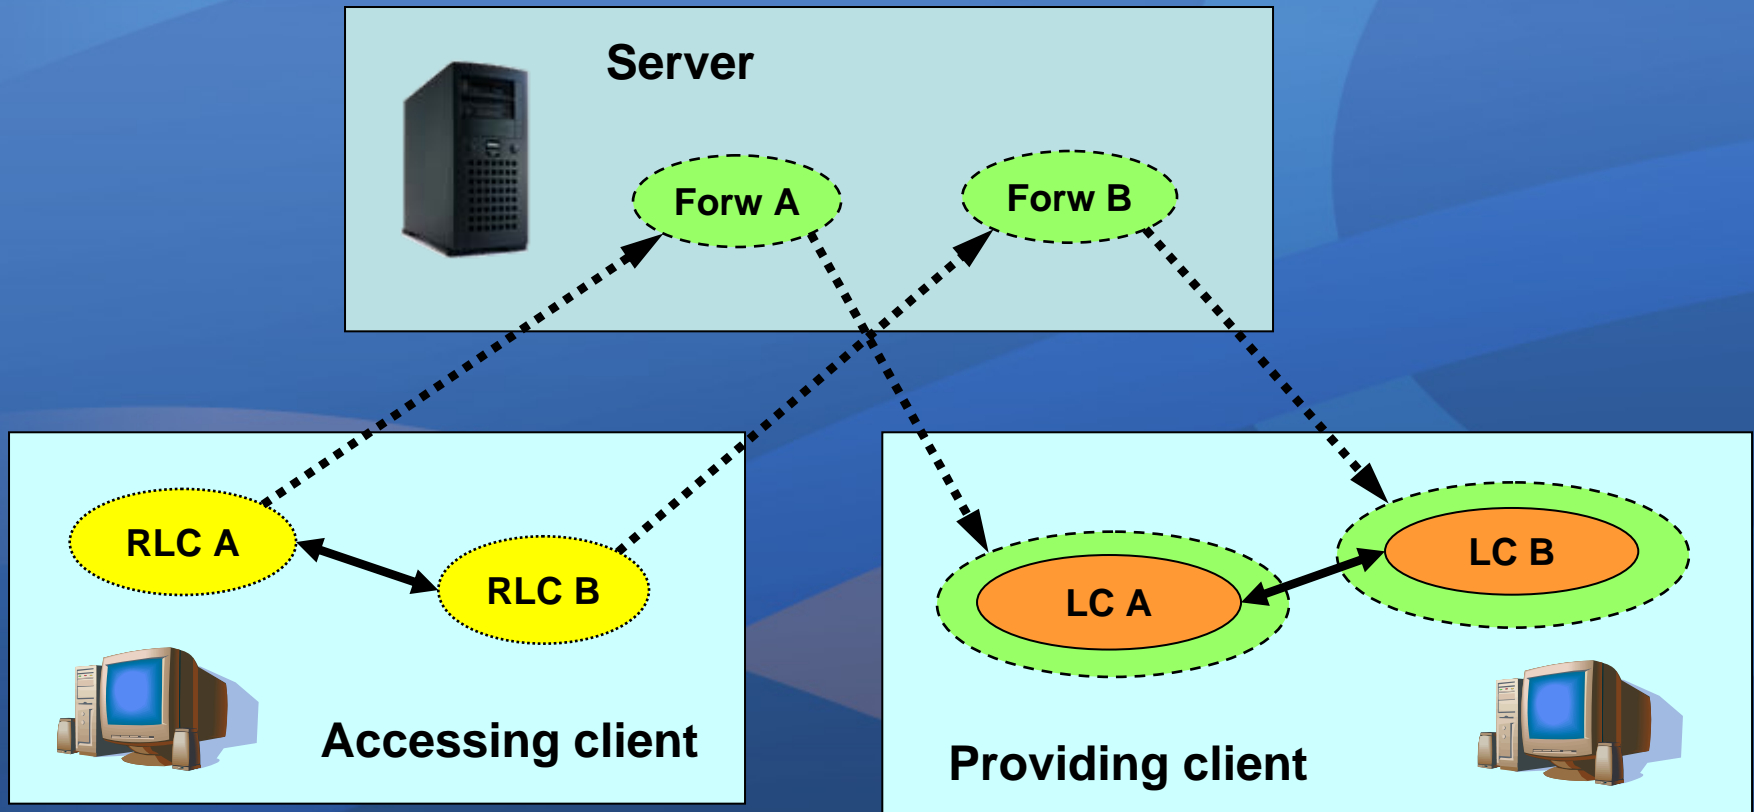


- Clients may be behind NAT or firewall
- Central register of clients and provided models
- Permission checking




Direct model linkage

- Improves performance of locally linked models dramatically !
- Server-side direct model linkage



Caching and piggybacking

- OpenMI not suited for distributed environment
- Many network roundtrips needed
- Caching + piggybacking solves the problem
- Result: during the computation only single roundtrip per time step
- Events decrease the performance



- Communication protocol platform independent
- Java client (and server) may be developed
- Linkage of Java and .NET models
- Linkage of models in Windows, Linux, Unix, ...
- Different computers (TCP/IP) / same computer (named pipes)



- Possibility to include Distributed OpenMI into official “OATC” release
- Technical improvements:
 - Communication failover, compression, encryption
 - Permissions
- True parallel computing ?