



Two-Level Load Balancing for



Ł. Opioła, M. Wrzeszcz, Ł. Dutka, R. Słota, J. Kitowski

ACC Cyfronet AGH
Department of Computer Science, AGH - UST

KUKDM 2015
Zakopane, Poland, March 11-13, 2015

Agenda



- Global data access
- onedata – a tool for global data access
- Addressed problem
- Two-level load balancing
- Test results
- Conclusions

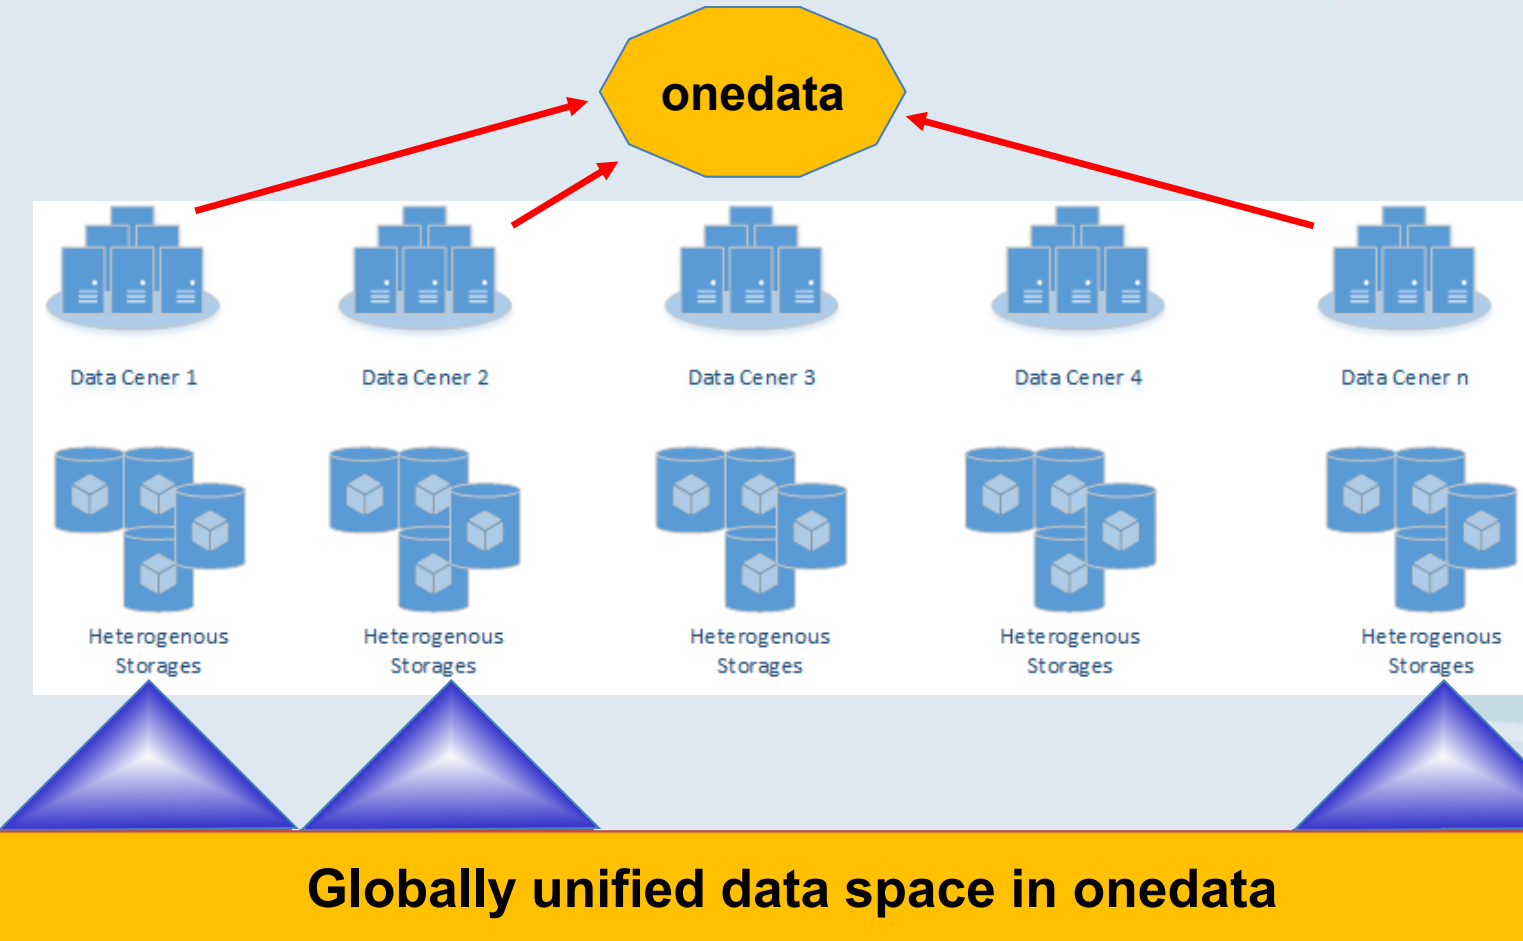


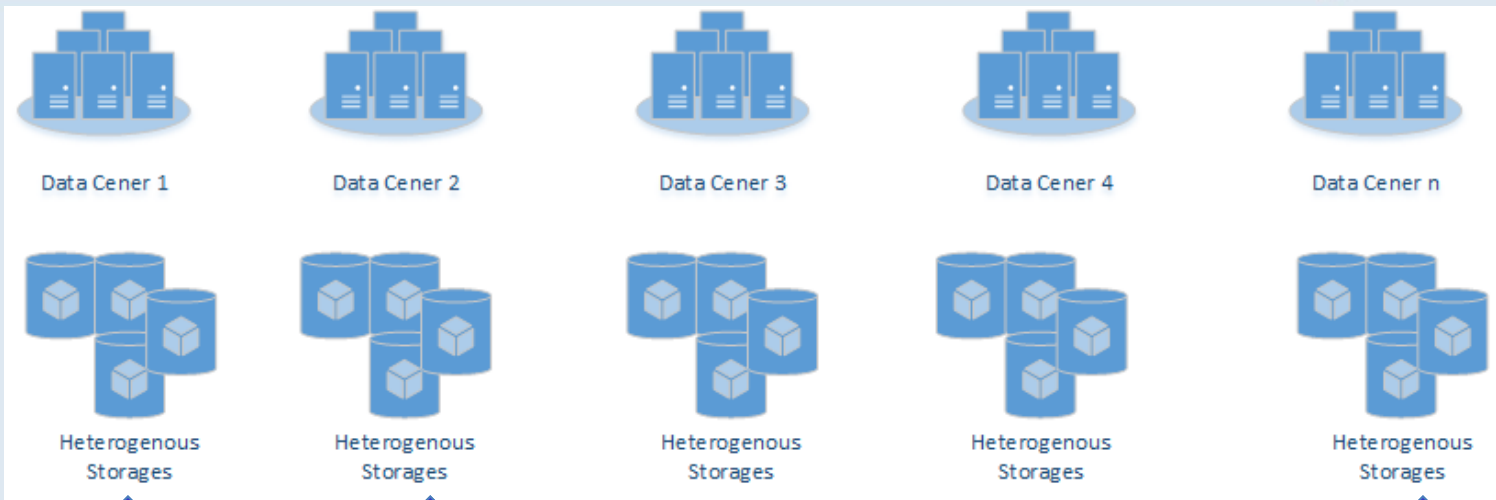
Global data access



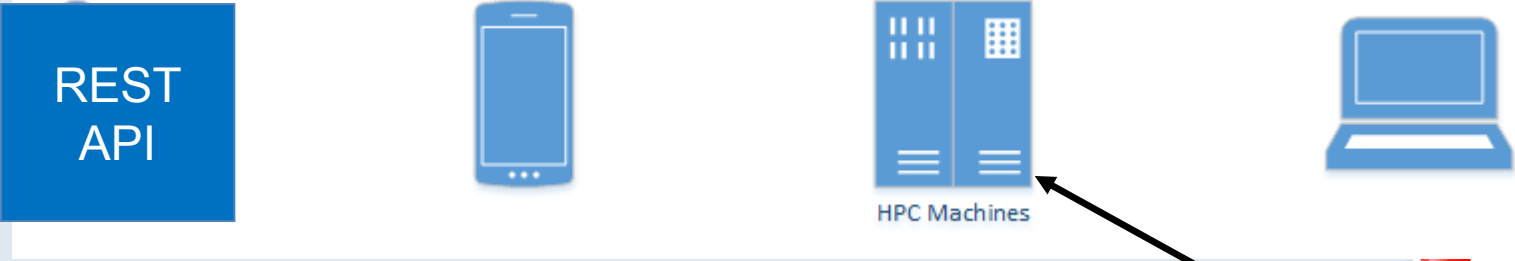
- Global, scientific data access in e-science
- User requirements
- Onedata – a powerful data globalization system







Globally unified data space in onedata



onedata

Efficient Access to data without need to migrate it before procesing



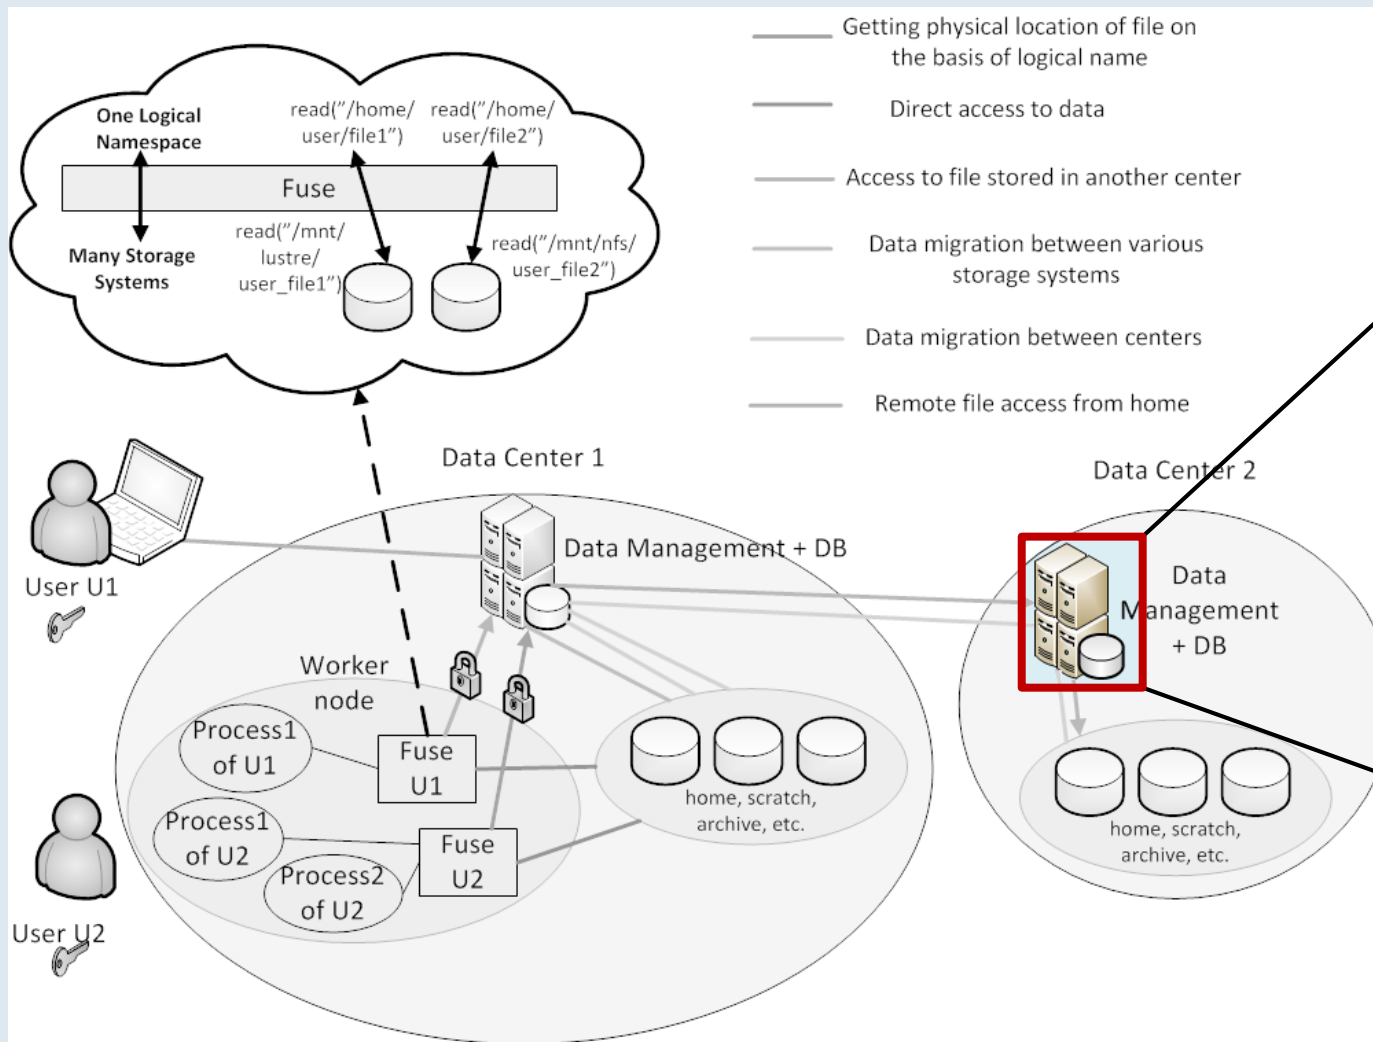
Straightforward support for groups and file sharing

The screenshot shows the onedata.org file manager interface. The browser address bar displays 'https://onedata.org/'. The page title is 'File manager' and 'Shared files'. The user 'Łukasz Opiota' is logged in. The current view is 'spaces'. A search bar is present. Below the navigation bar is a table of shared files.


Name	Size	Access	Modification
..		2014-09-01 17:33:32	2014-09-01 17:33:32
eksperyment 1		2015-03-08 20:17:28	2015-03-08 20:17:28
wspdzielone		2015-03-08 20:18:28	2015-03-08 20:18:28
zadanie 3		2015-03-08 20:19:29	2015-03-08 20:19:29



onedata architecture



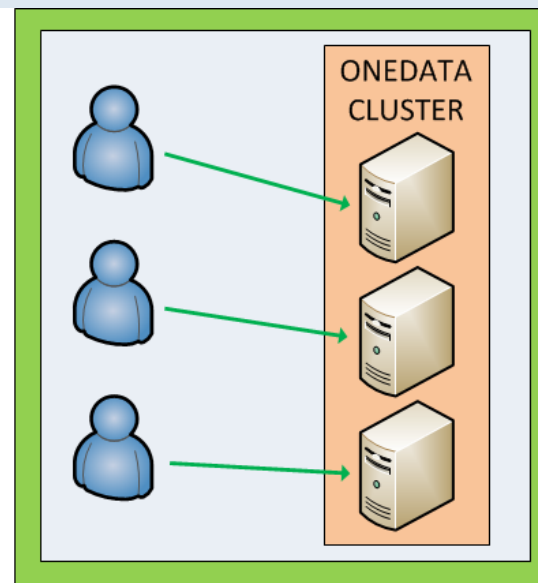

- Getting physical location of file on the basis of logical name
- Direct access to data
- Access to file stored in another center
- Data migration between various storage systems
- Data migration between centers
- Remote file access from home



Addressed problem


Effective load balancing in a cluster of nodes:

- Scalability
- Efficient resource utilization
- High availability
- Bottlenecks avoidance

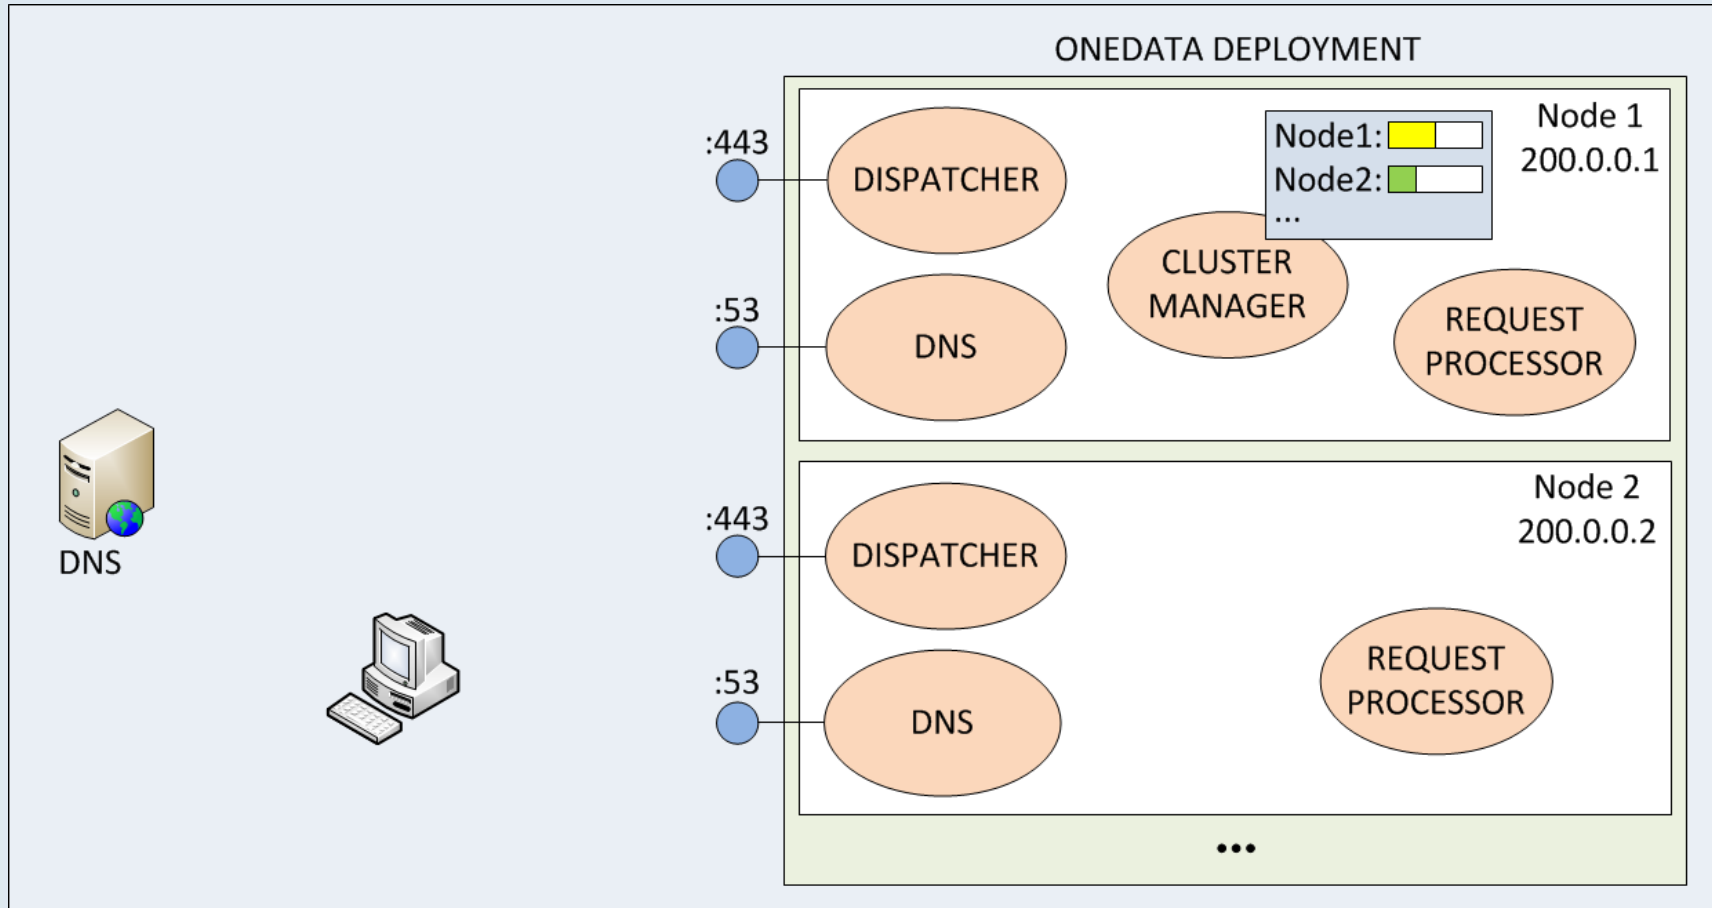


Two-level load balancing

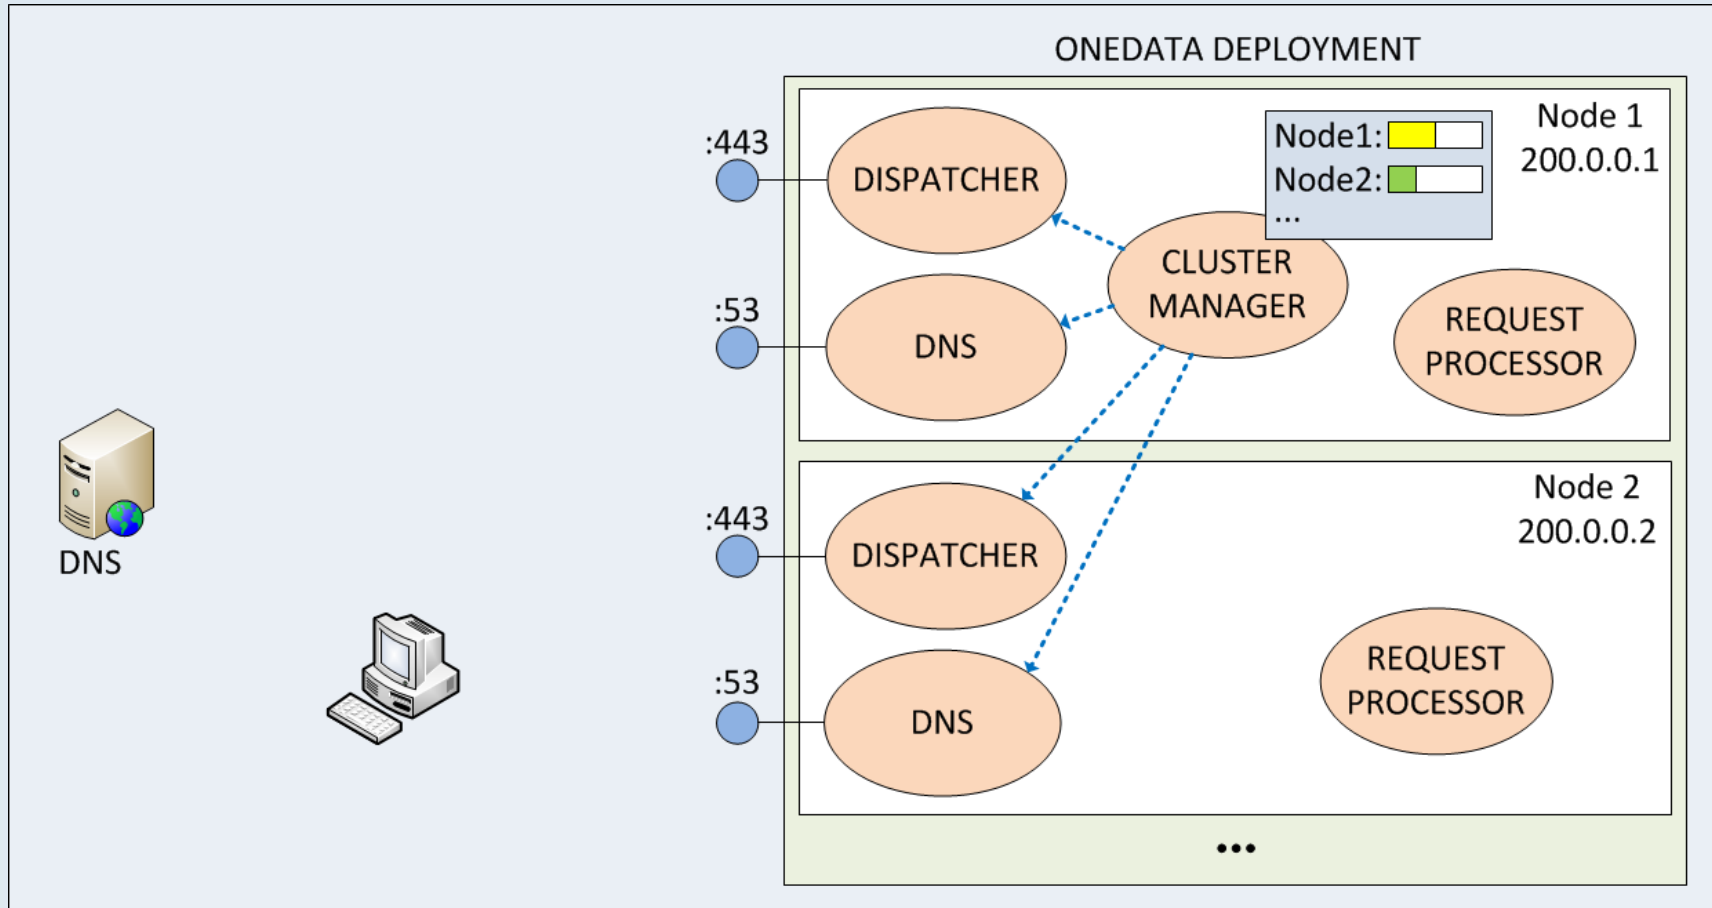
- Level I: DNS
 - Integrated with the cluster
 - Distributes load among interfaces
 - Elasticity in case of node overload or failure
- Level II: Request dispatcher
 - Reacts to load fluctuations
 - Handles edge cases
- Is either one not enough?



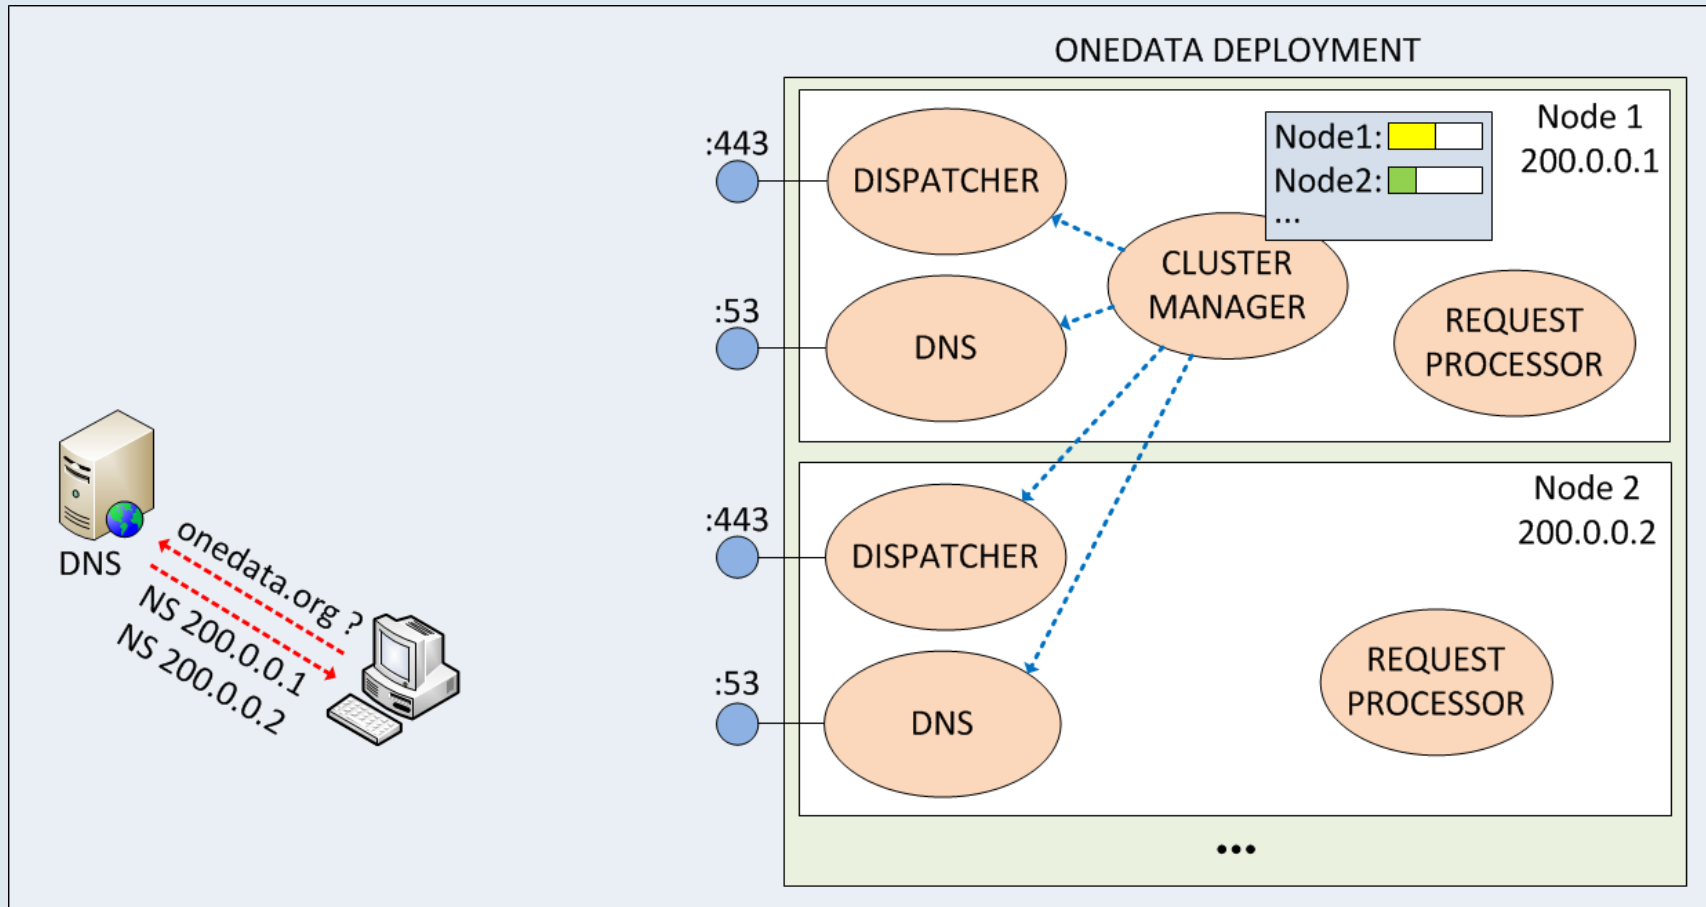
Two-level LB at work



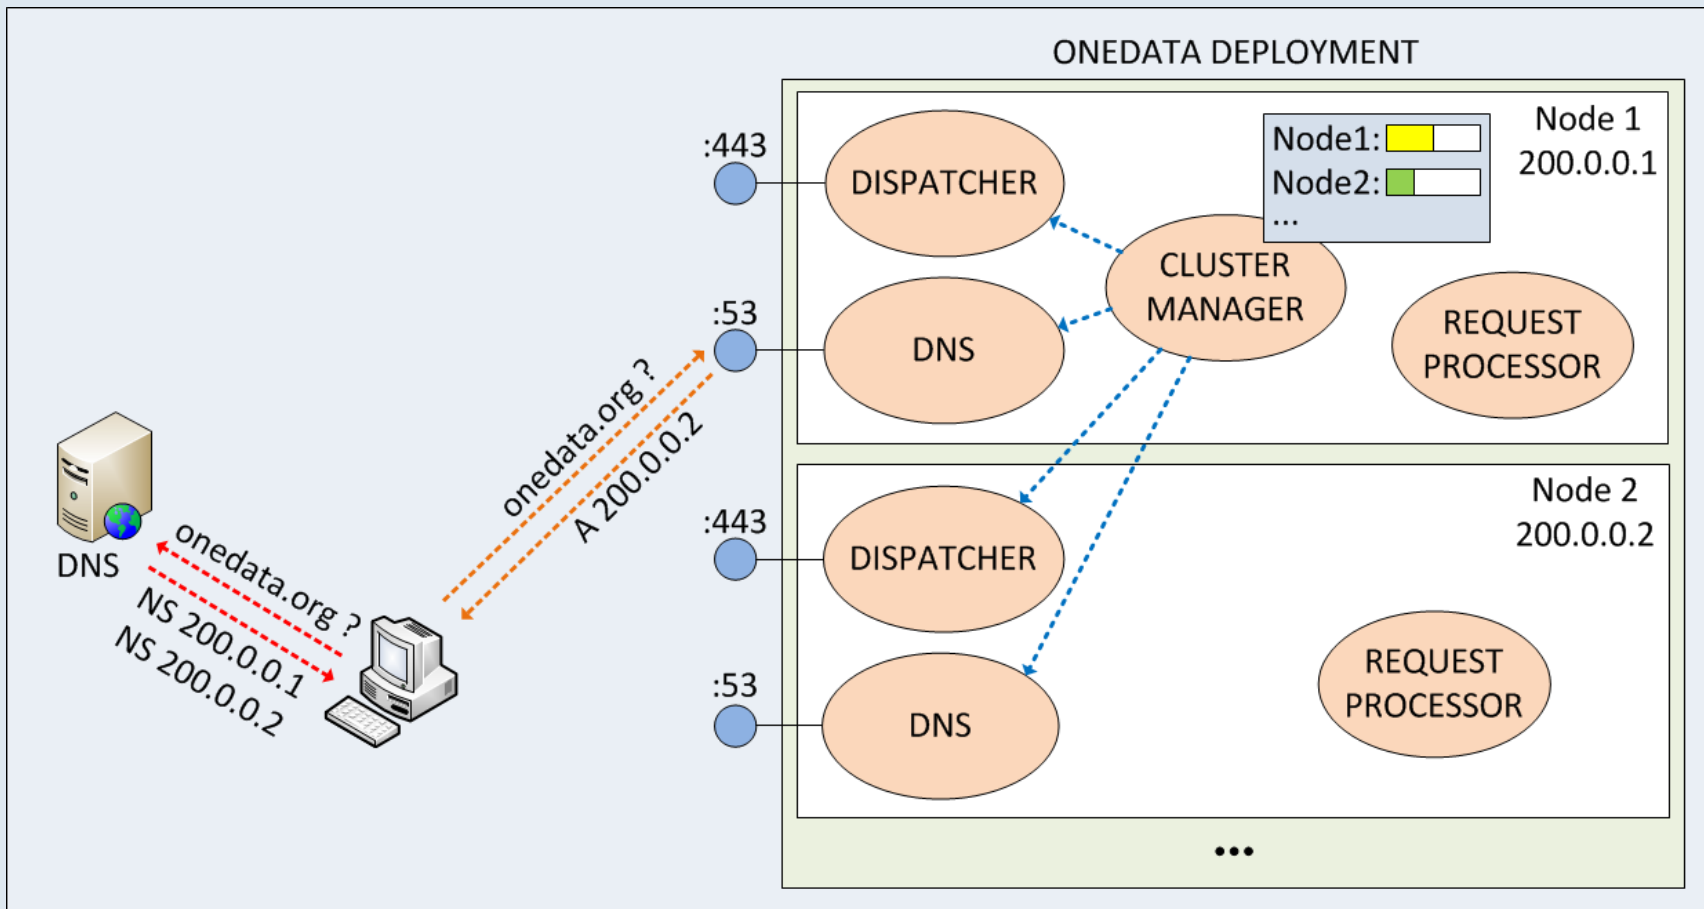
Two-level LB at work



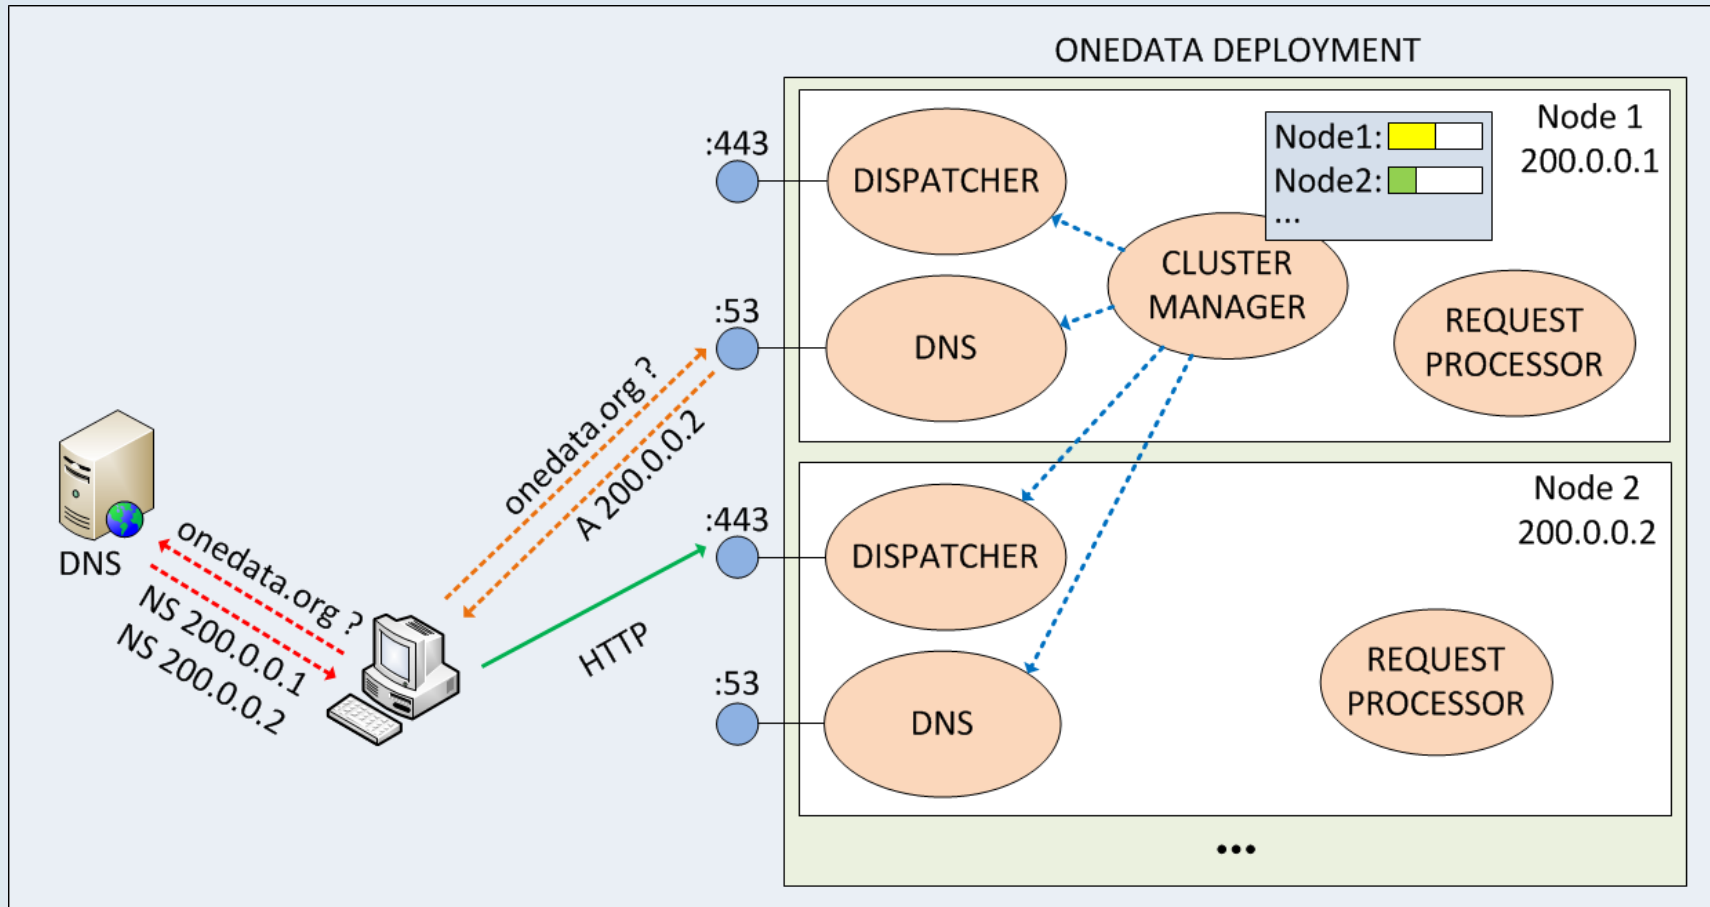
Two-level LB at work



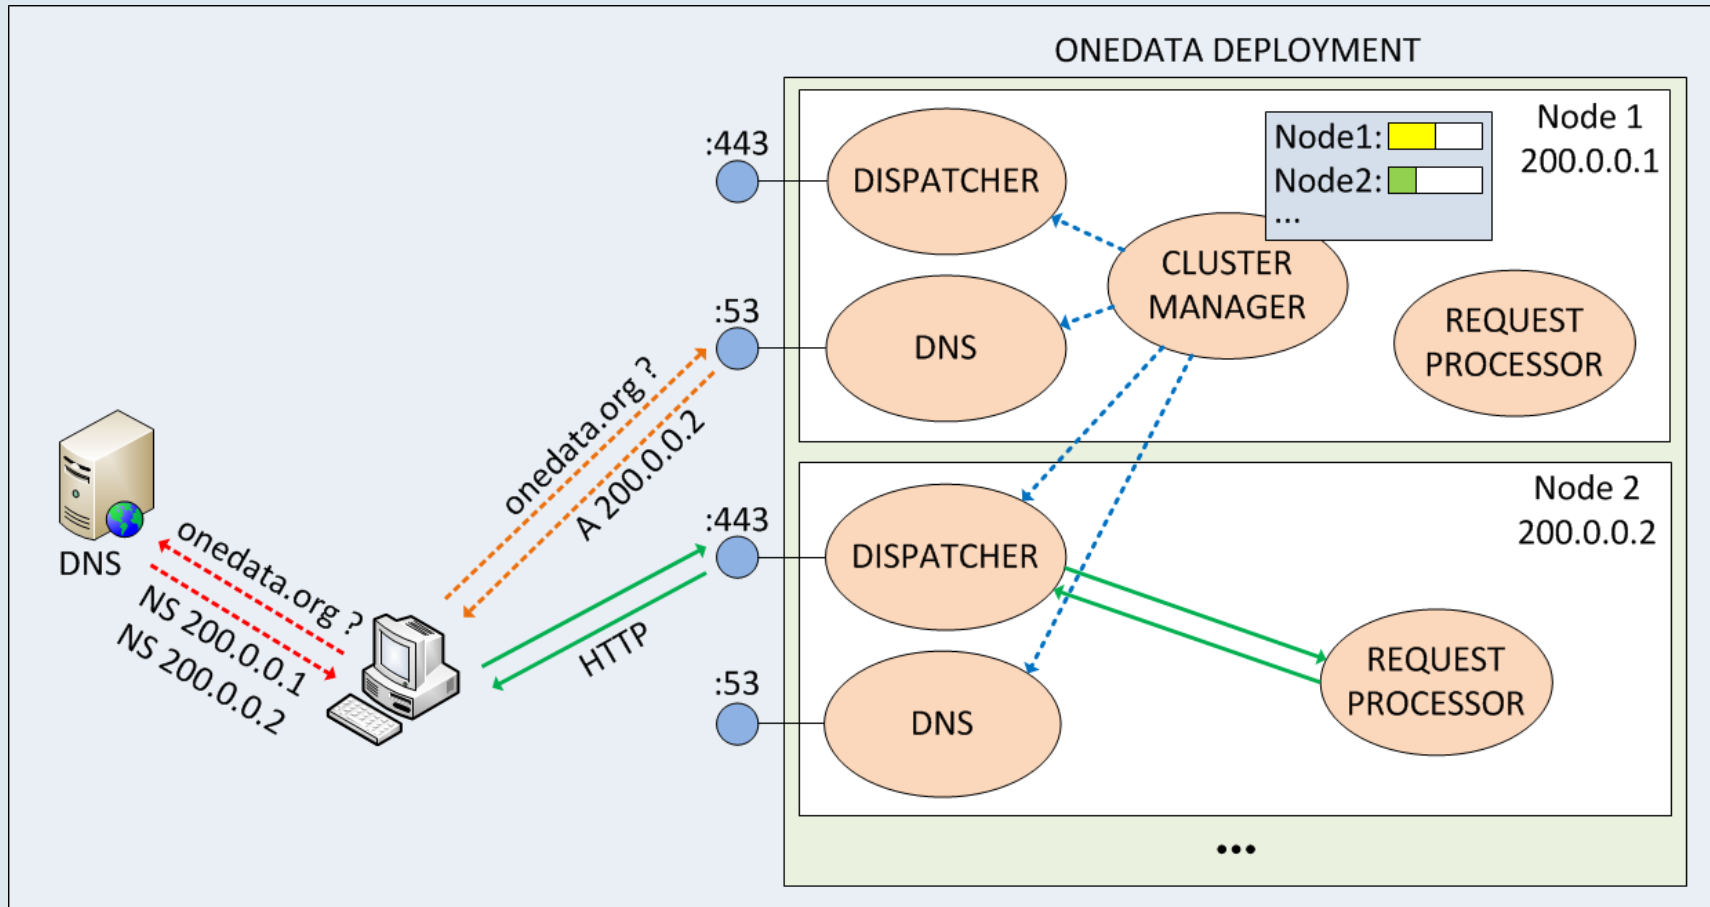
Two-level LB at work



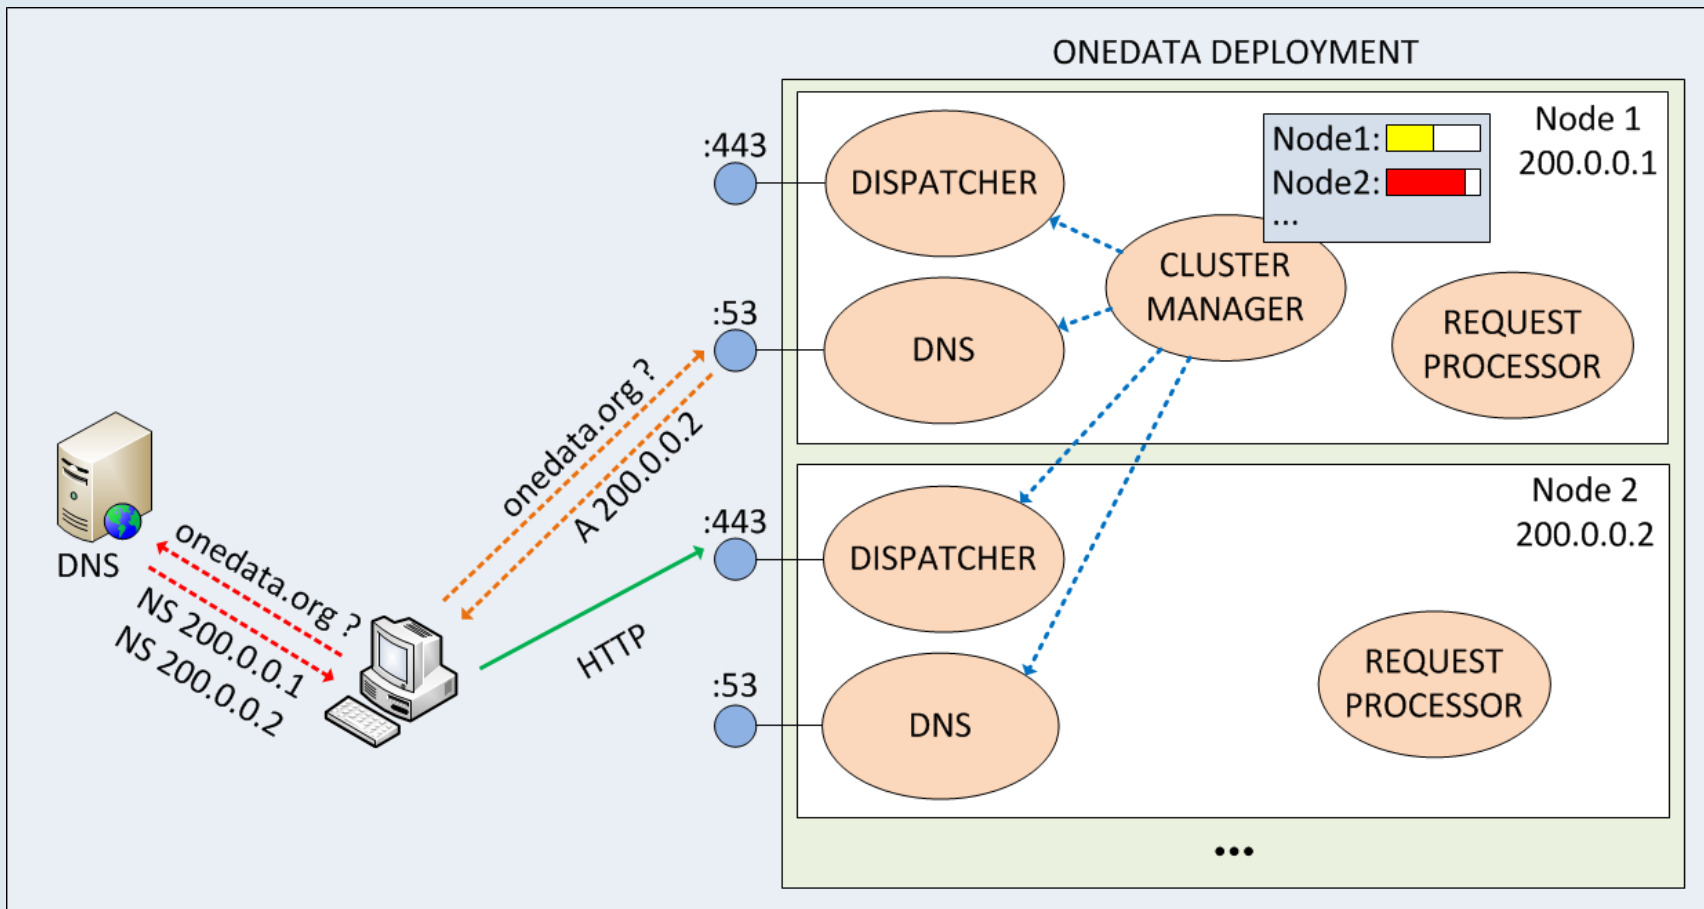
Two-level LB at work



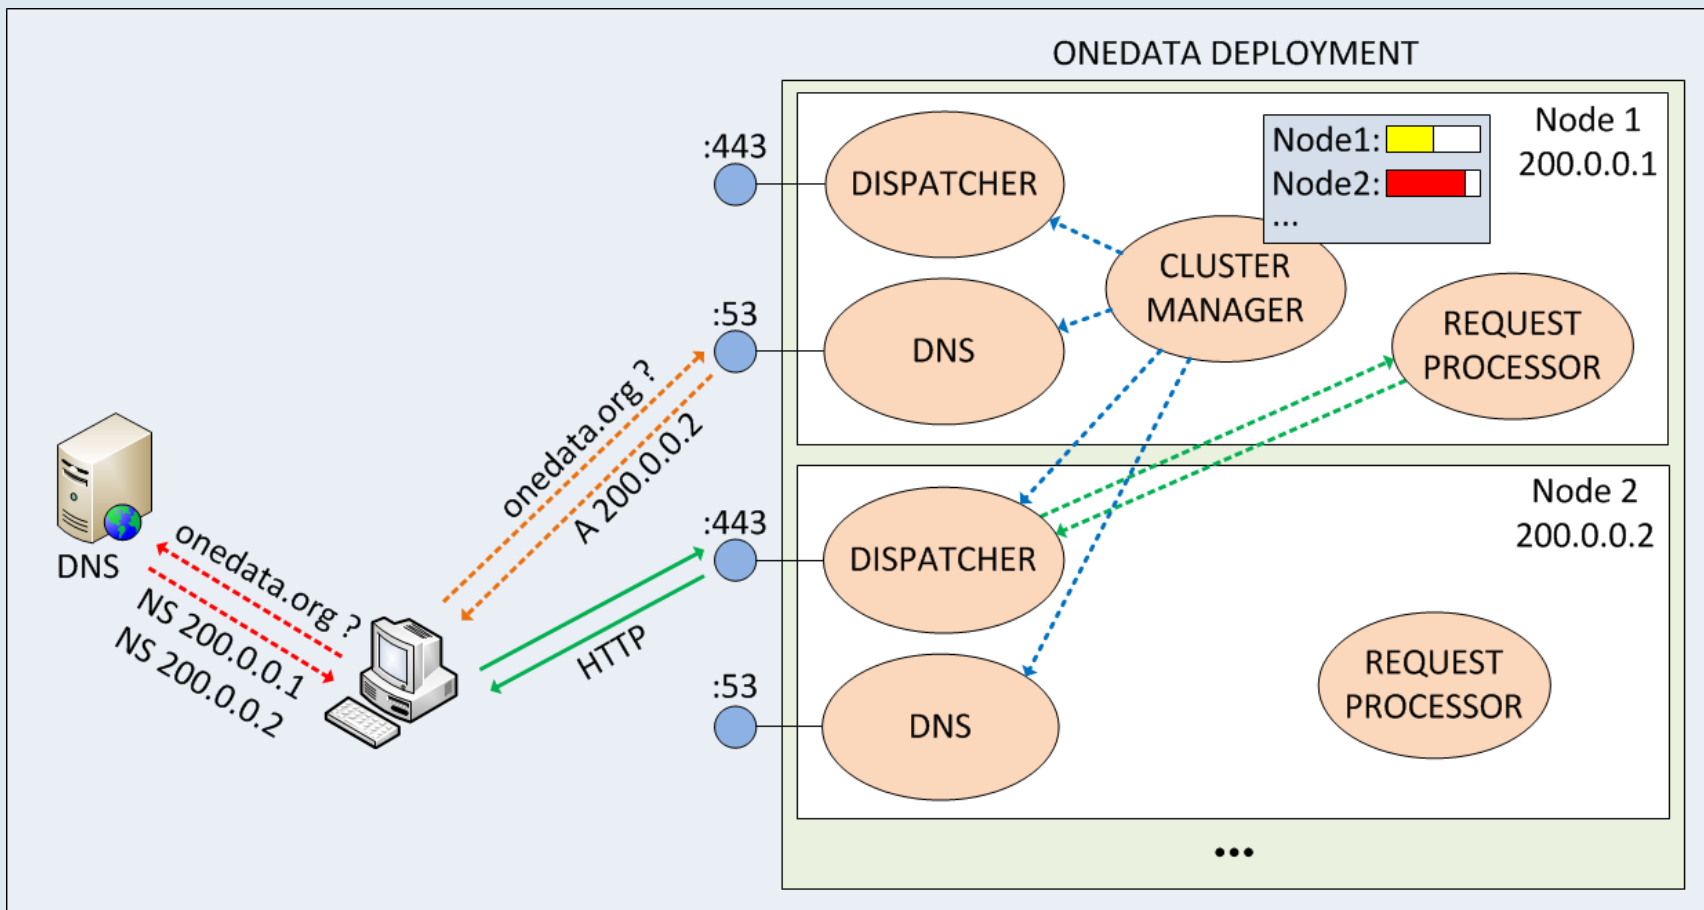
Two-level LB at work



Two-level LB at work



Two-level LB at work




Test setup

Requests:

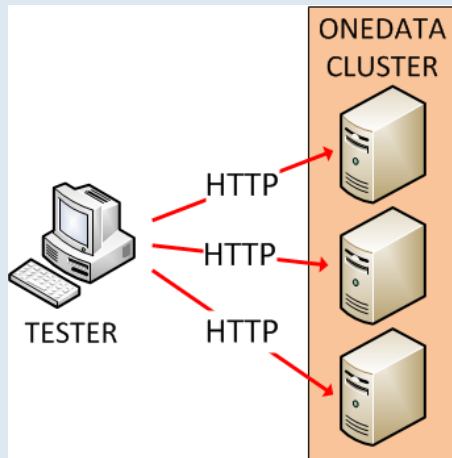
- **short** (quick responses)
- **long** (computationally requiring)
- **heavy** (lots of network I/O)

Test cases:

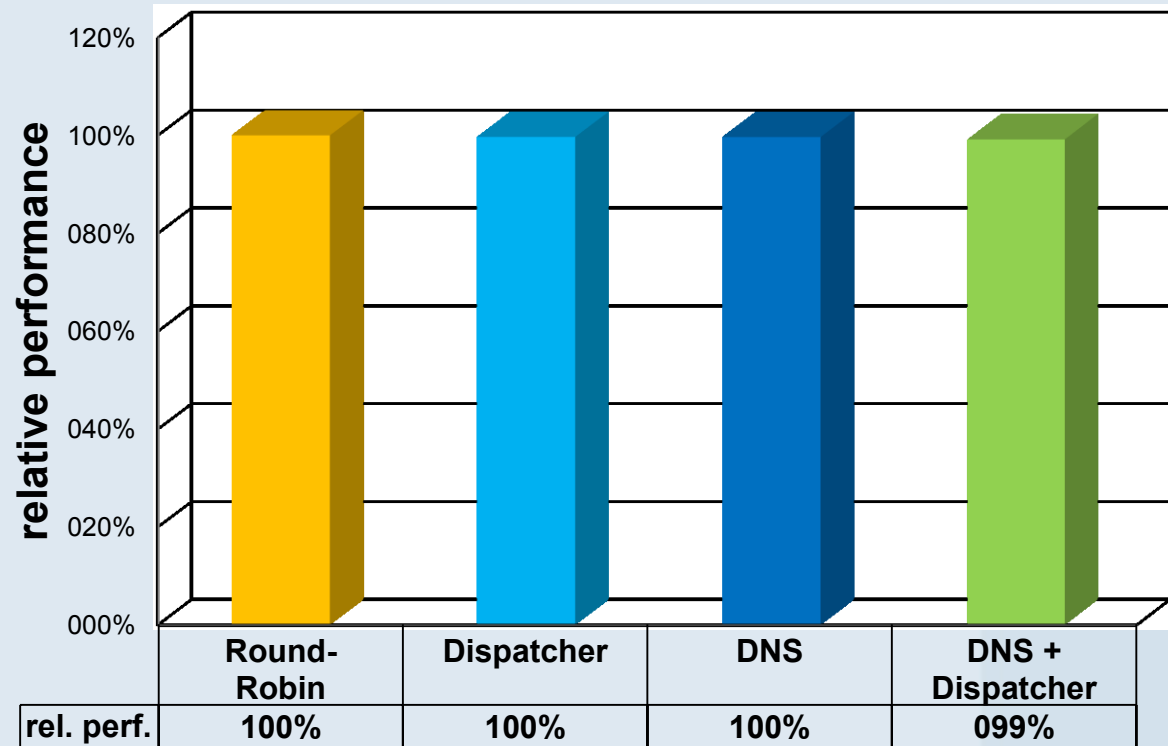
- Round-Robin (for reference)
- DNS
- Dispatcher
- DNS + Dispatcher




Test results – long & short & heavy



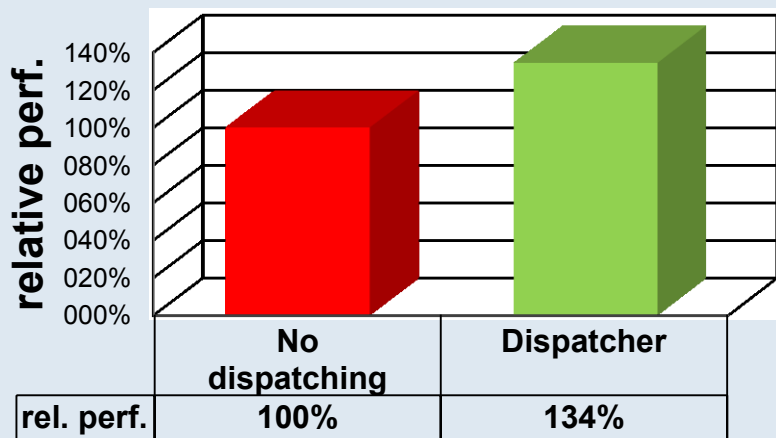
Throughput – average for all request types



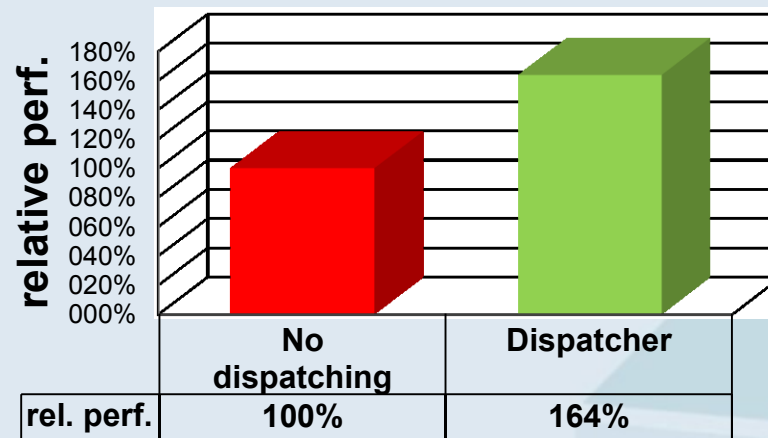
Test results - dispatcher



long requests



very long requests



Conclusions



- Two-level load balancing was successfully implemented and has proven a viable solution
- It combines best features of the two approaches
- The overheads do not have impact on the performance of the whole system
- onedata can rise up to challenges of global data access





Thank you

onedata homepage: <http://www.onedata.org>

VeilFS in PLGrid: <http://www.veilfs.plgrid.pl>