The Information System

A Status Report

Laurence Field

www.eu-egee.org

Information Society and Media

Introduction

- **Brief Overview of the information system**
- **Observed problems**
- **Explanation of the problems**
- **Investigations undertaken**
- **Preliminary results**
  - Stability
  - Query rates
  - Query scaling
  - Loading Effects
- **Short term solutions**
- **Long term solutions**

The Information System

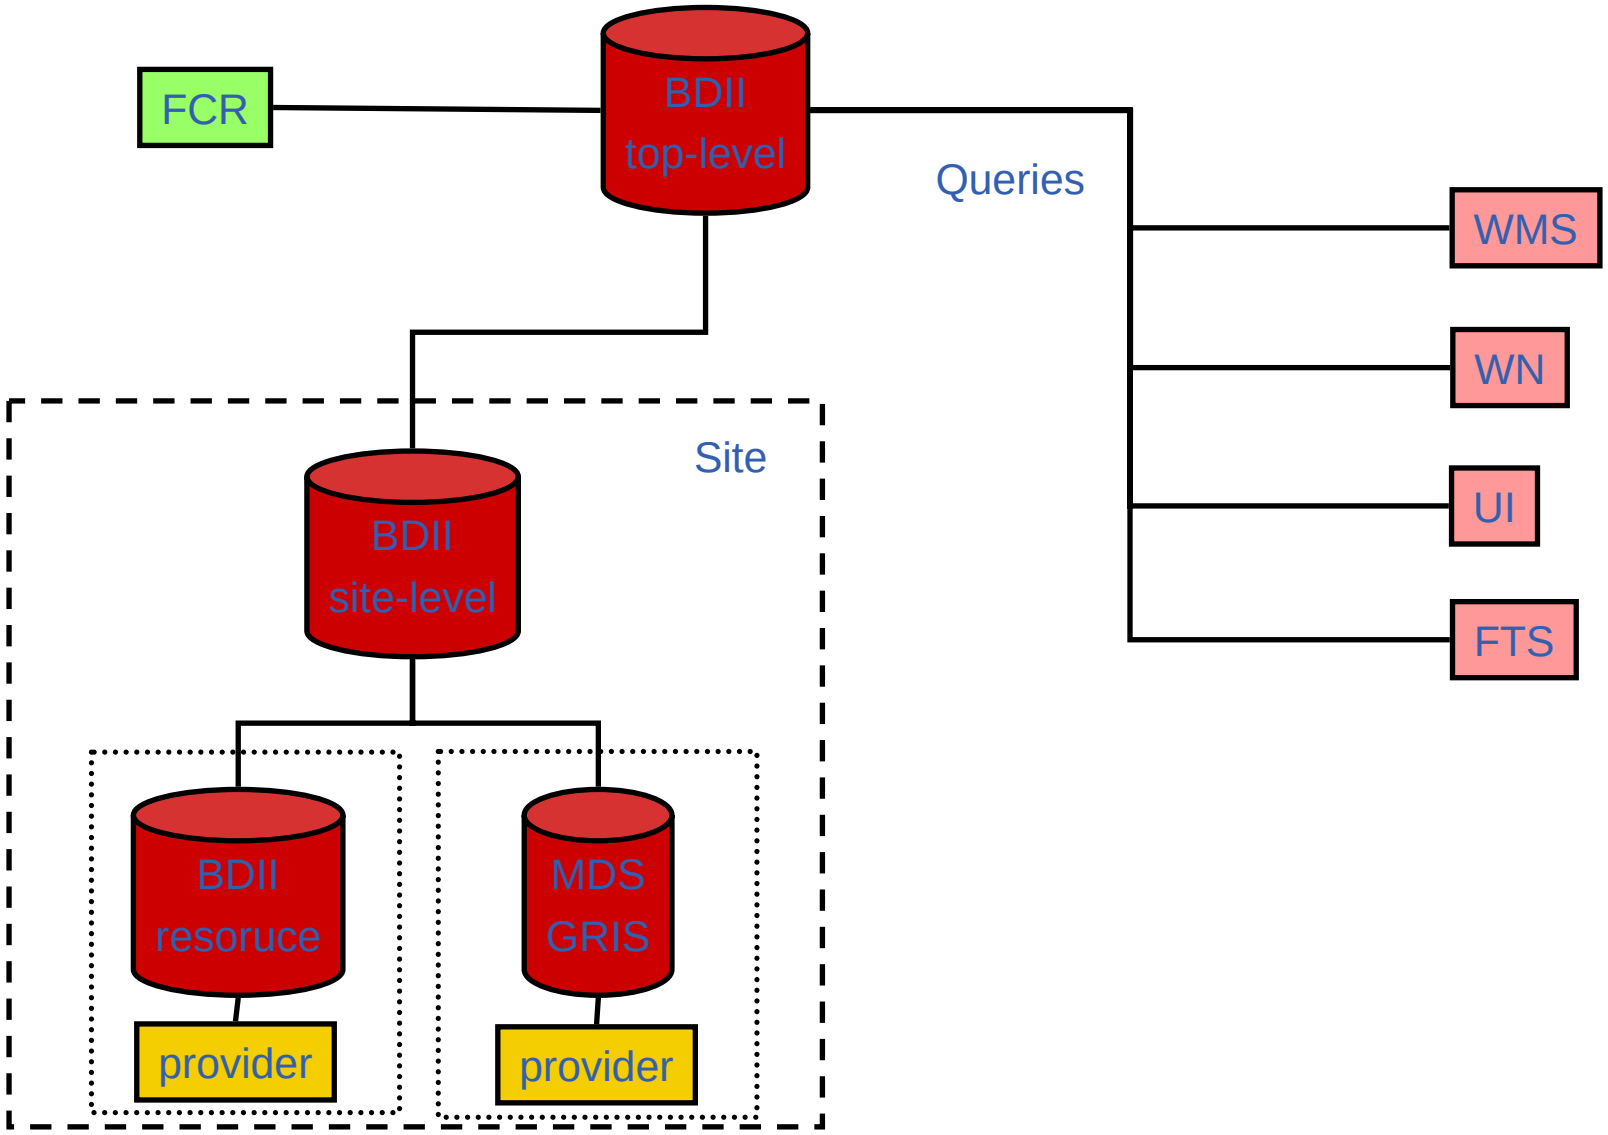

Inside A BDII

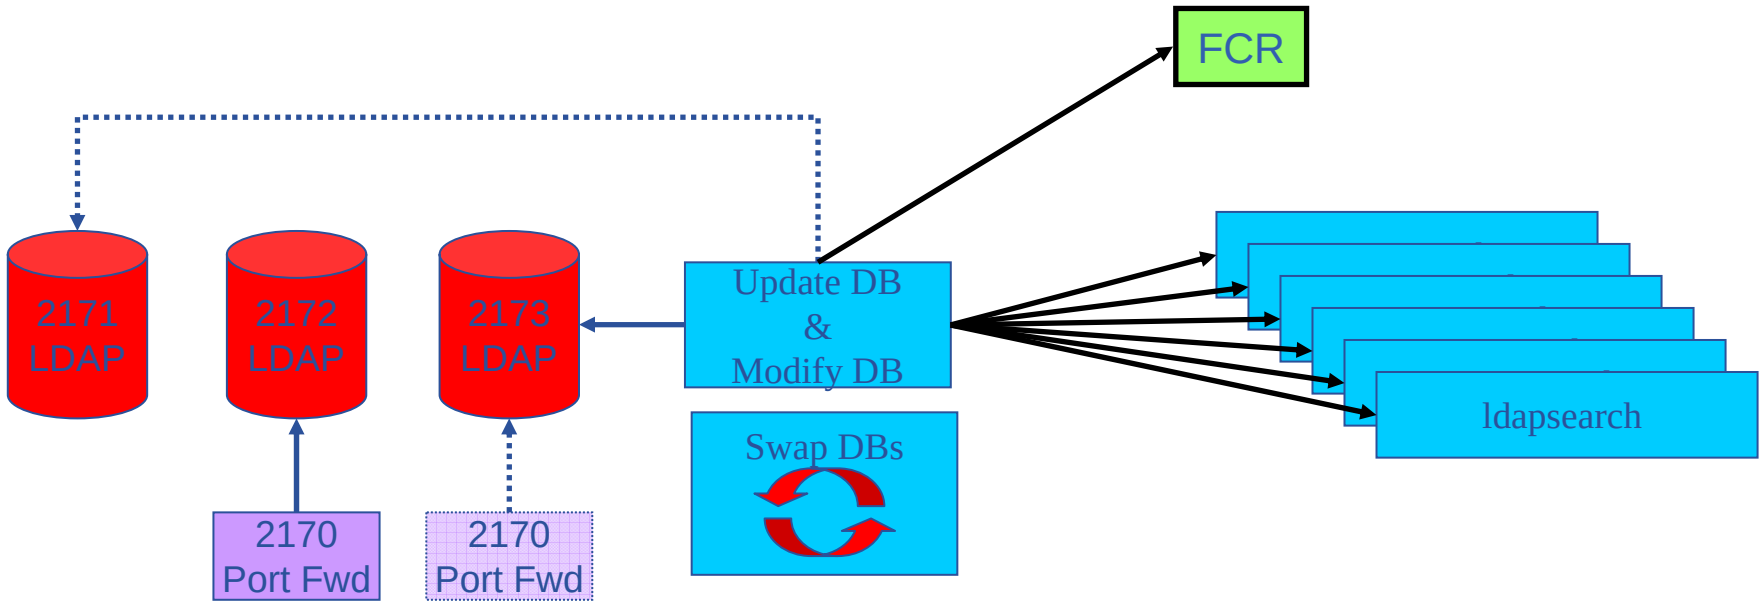

Load Balanced BDII

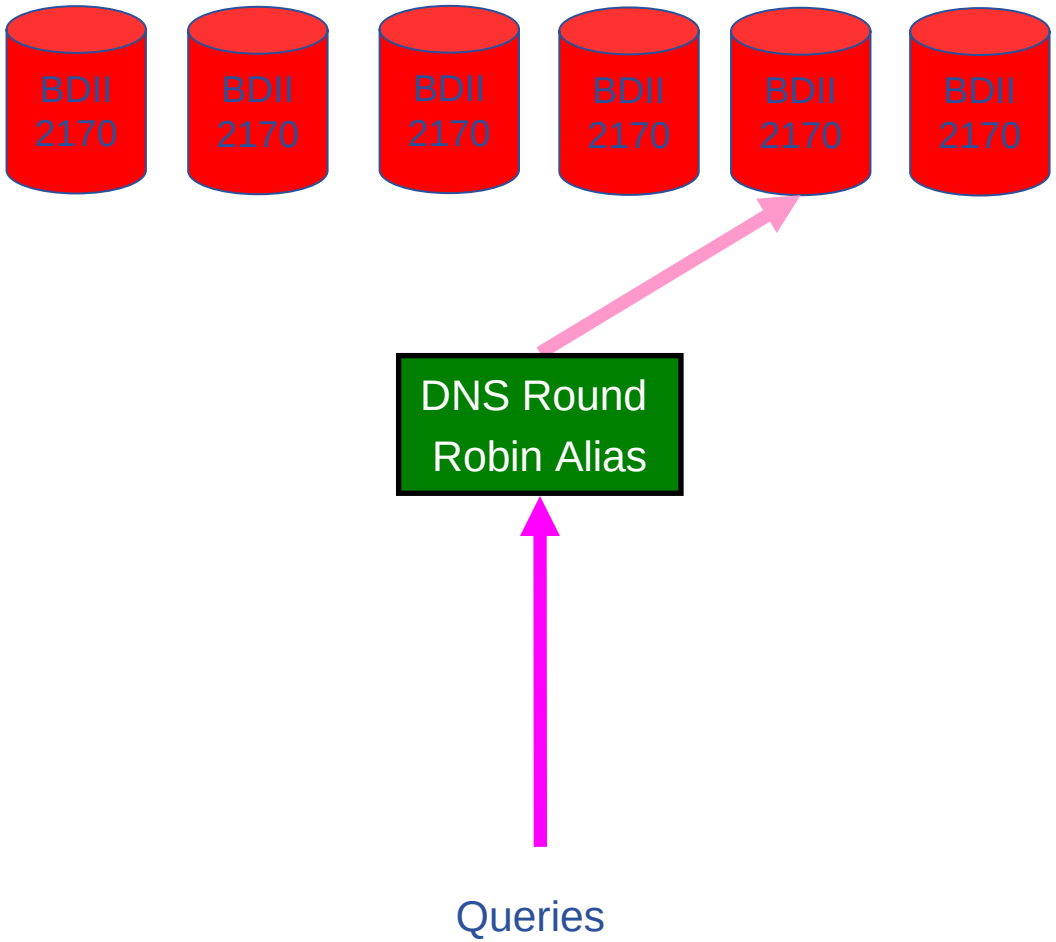

Observations

- **Queries to the top-level BDII time-out**
  - 15s for lcg-utils
- **Some site information is unstable**
  - Dropping in and out of the top-level BDII
    - Affects FTS pre-scheduling
- **Some CE information is unstable**
  - Dropping in and out of the top-level-BDII

- **Problems observed in information system**
  - not always due to information system
    - It is just where the problem is visible
  - Many problems at the information providers level
    - Due to either poor configuration
    - Poor fabric management affecting information providers

Problems

- **Slapd is adversely affected by machine loading**
  - High load can cause the slapd to not respond
    - High load can also affect the information providers
- **This is the main cause of CE disappearing**
  - High load due to many jobs
- **Same reason why Site BDII disappearing**
  - Usually due to the CE and BDII co-hosted on the same node

- **Query rate**
  - Querying a single point is creating huge query rates
    - 40K+ WokerNodes trying a DNS attack ☺
  - Assuming that all information is highly dynamic
    - Freshness of 30s
  - Most information required is static
    - Freshness of hours is sufficient

Investigations

- **Monitoring top-level BDII for**
  - Information instabilities
  - Time-outs
  - Query rate
  - Query types

- **Scalability tests on the slapd**
  - Effect of parallel queries
  - Effect of data size
  - Affect of loading

Monitoring Information Stability

- **Top-level BDII monitoring**
  - Measurements made every 30s
- **Time-outs of 10 seconds**
  - On average 2 per day
    - Occasionally 1 hour of instability with more than 2 time-outs
- **Site information disappearing**
  - Site BDII failing to respond
    - 0-100 times per week
      - *Dependent on site!*
    - Most due to the site-level bdii co-hosted with the CE
- **CE information disappearing**
  - High load causes slapd not to respond
    - 0-100 times per week
      - *Dependent on site!*

Queries To The Top-Level BDII

- **Analysis of a log file from one top-level BDII over 4 hours**
  - Multiple by 8 to get the values for the load balance service

No of Queries	Query
6075	Find the Close CE to an SE
5475	Find the VOs SA for an SE
5043	Find all SRMs
4791	Find an SE
2432	Find the Close SE to a CE
2117	Find all Services for a VO
664	Find all CEs for a VO
638	Find all SAs for a VO
479	Find all SubClusters
448	Find the GlueVOView for a CE

Queries To The Top-Level BDII

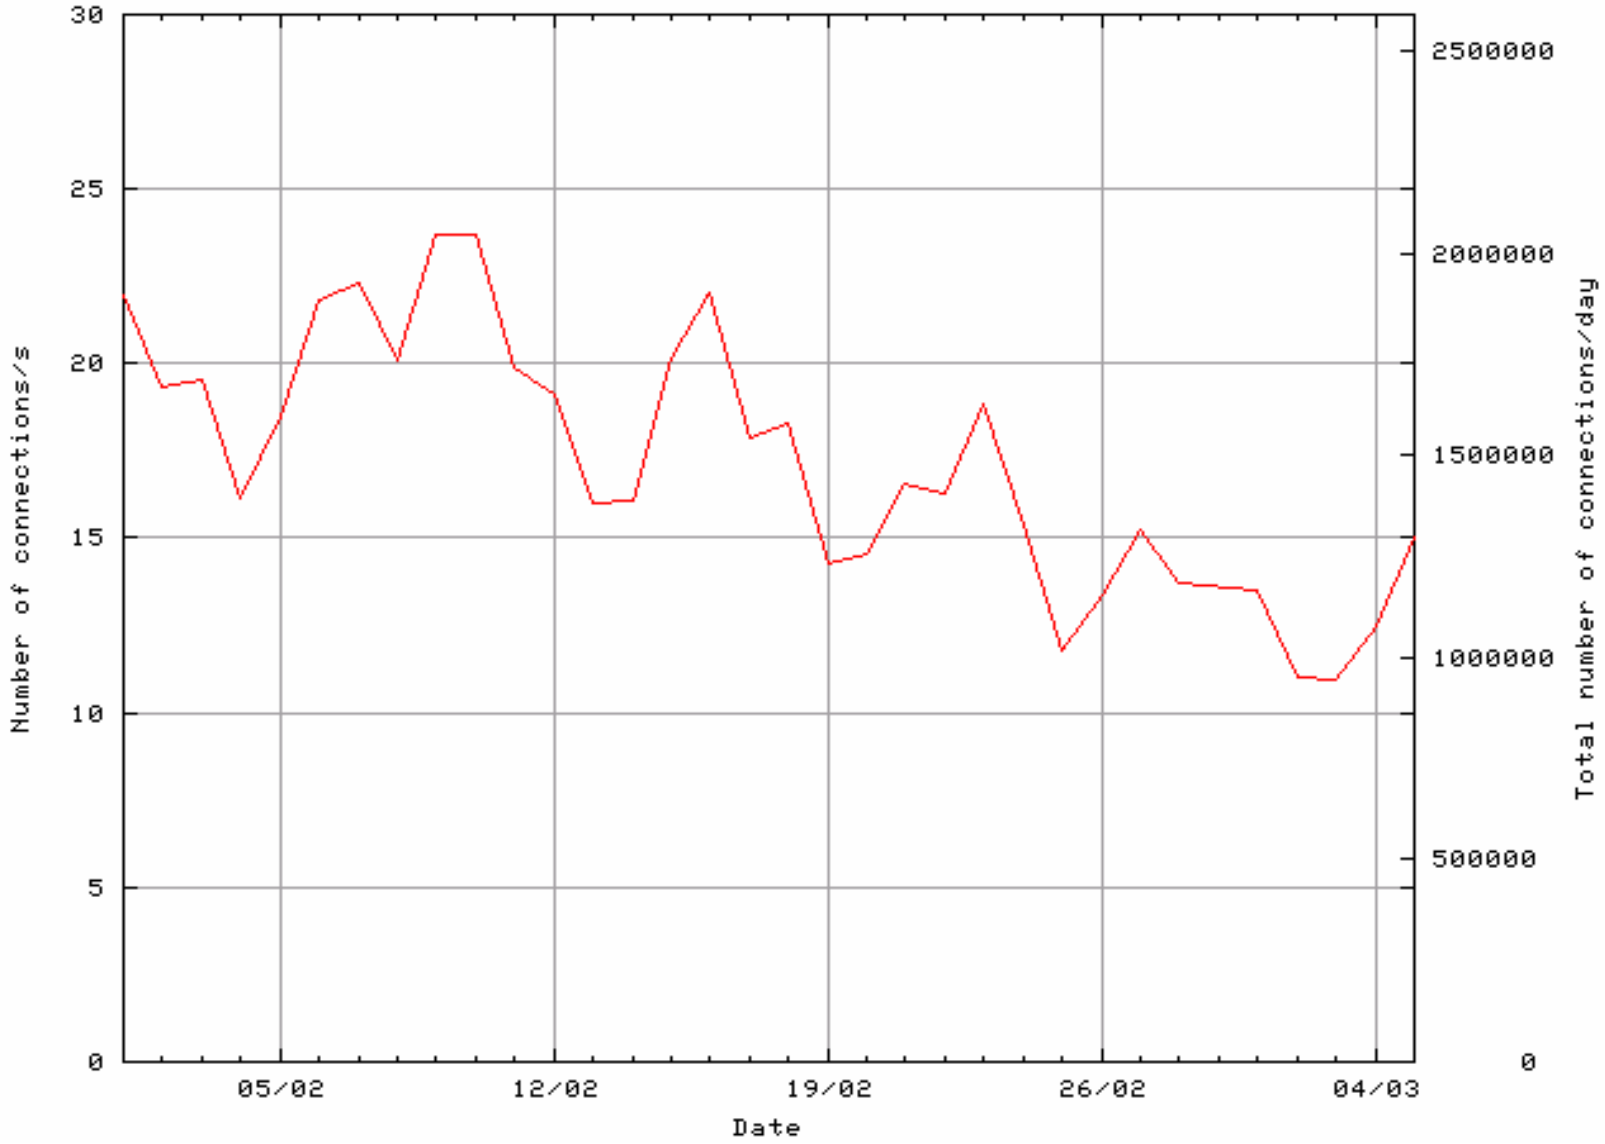

Query Rates By Host

Queries	Host
74828	rb101.cern.ch
40963	nat-outside-fzk.gridka.de
40552	lcgvm.triumf.ca
33370	rb114.cern.ch
31148	cert-rb-06.cnaf.infn.it
23461	rb123.cern.ch
20864	nat005.gla.scotgrid.ac.uk
80071	cclcglhcb.in2p3.fr

Query Scaling Of slapd

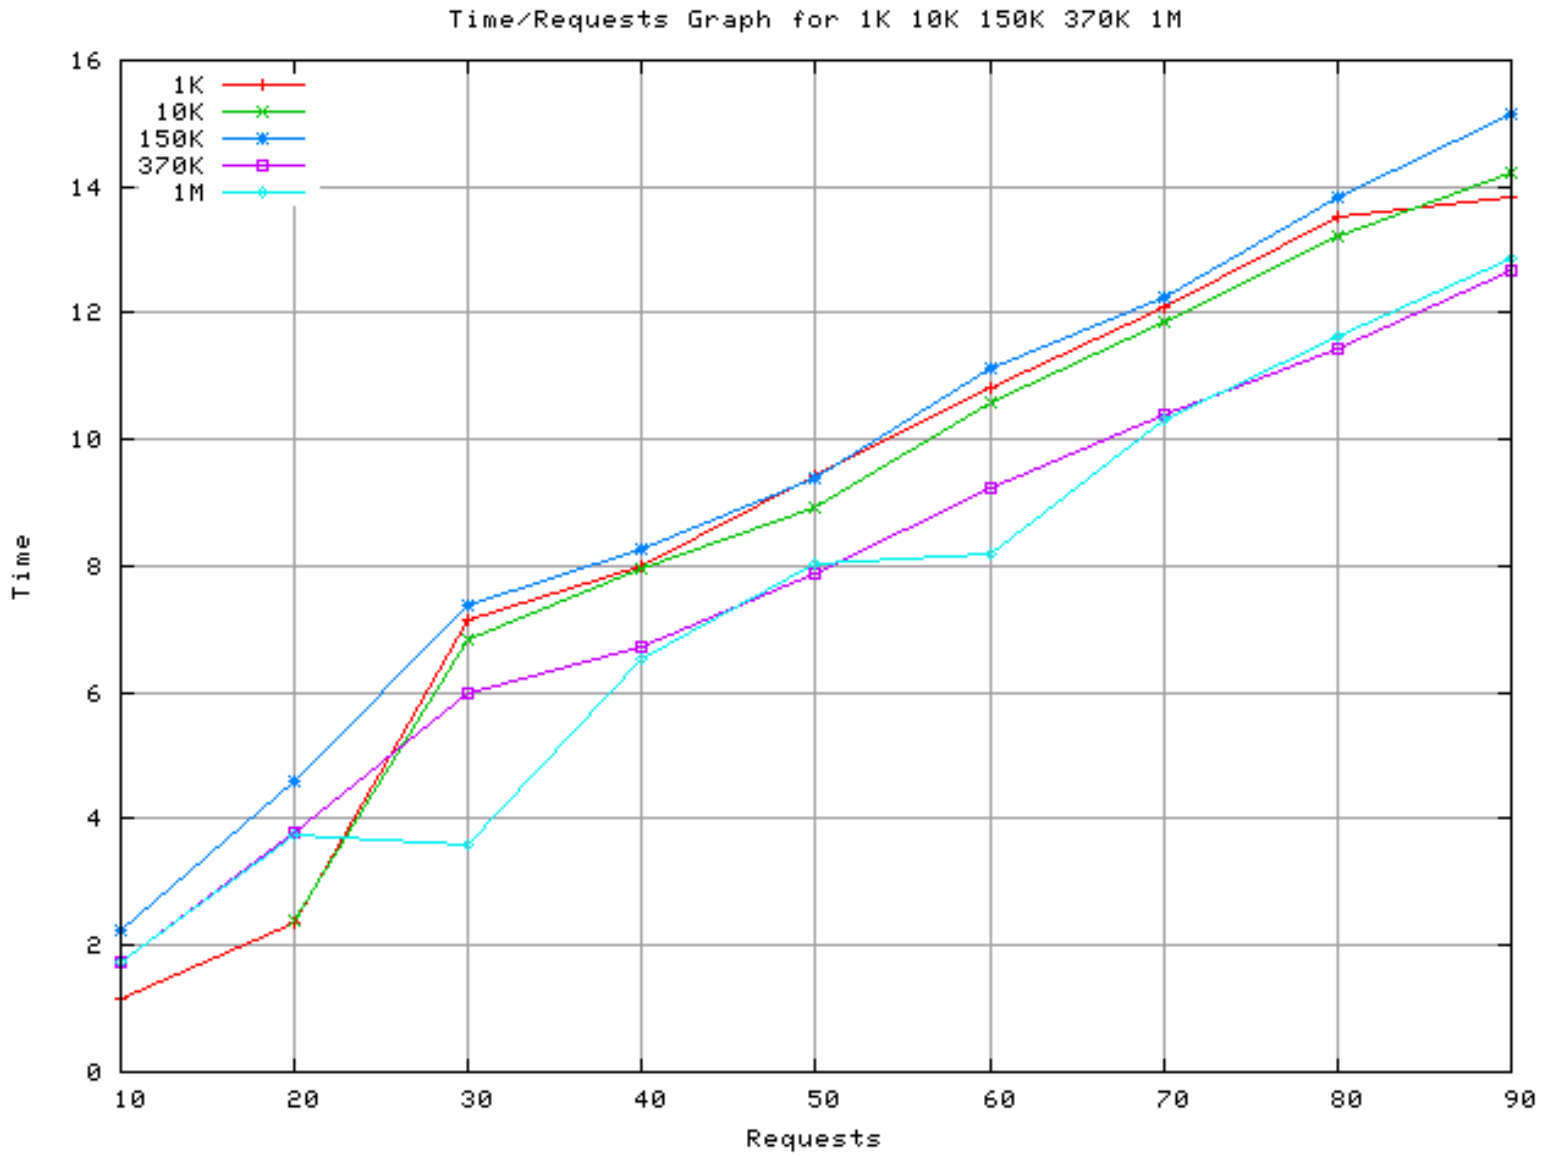

Query Scaling Of slapd

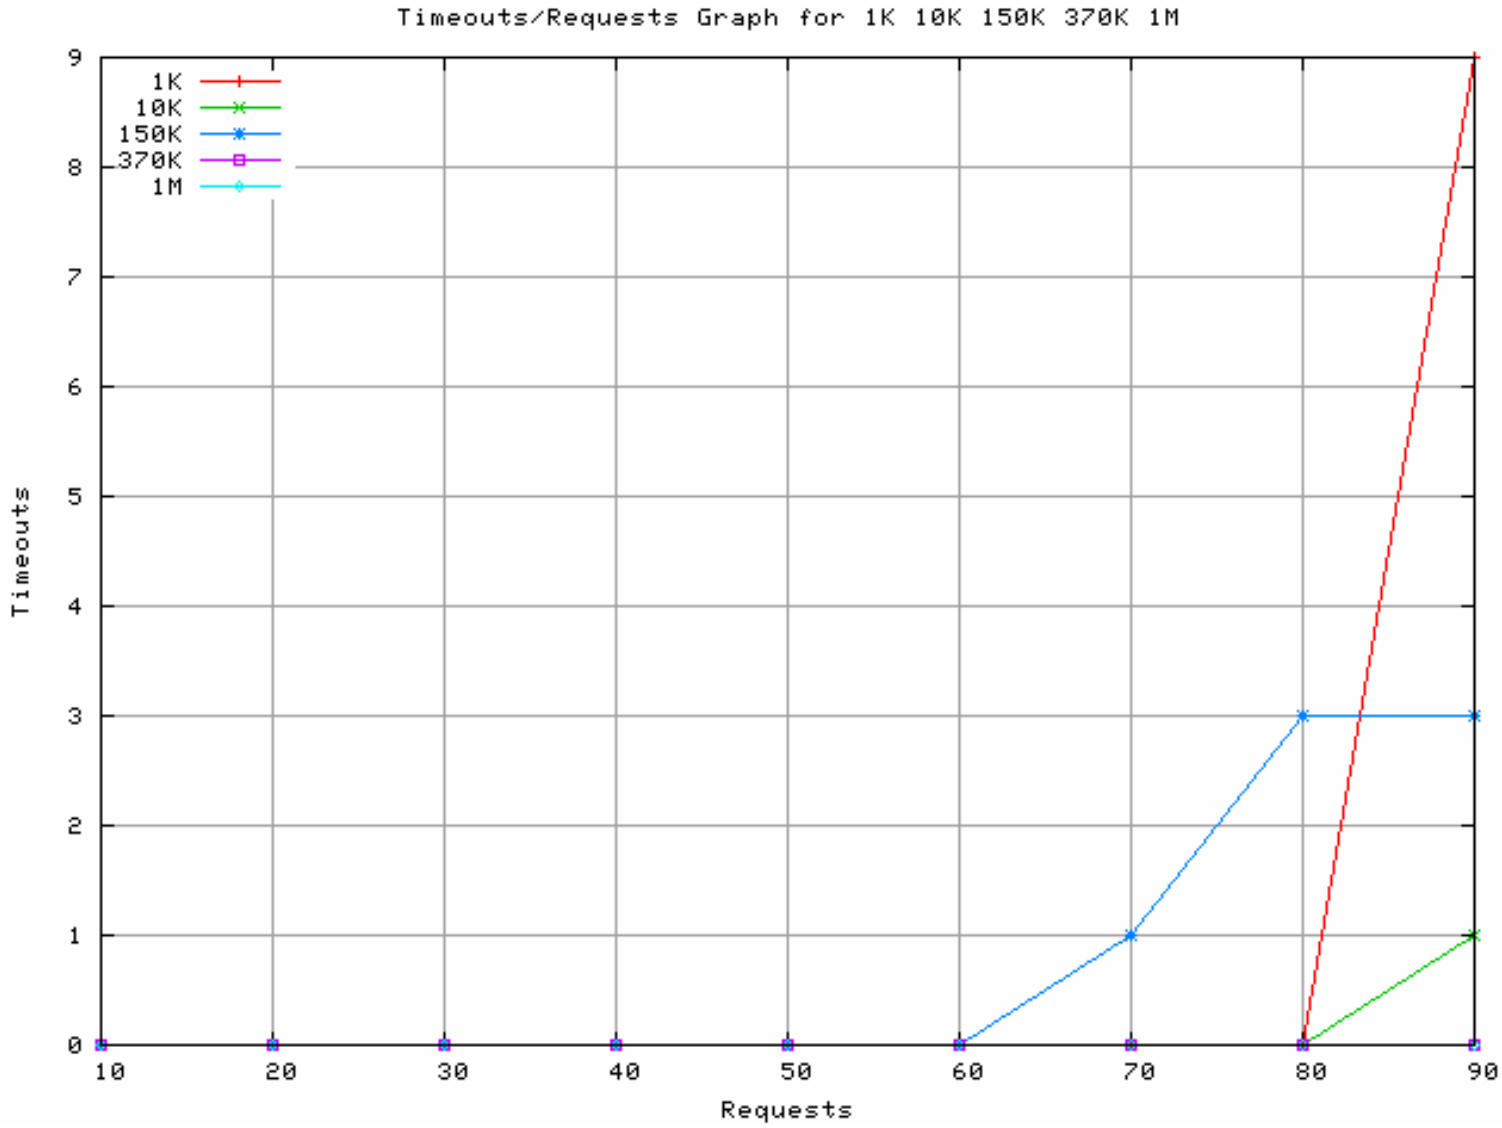

Query Scaling Of slapd

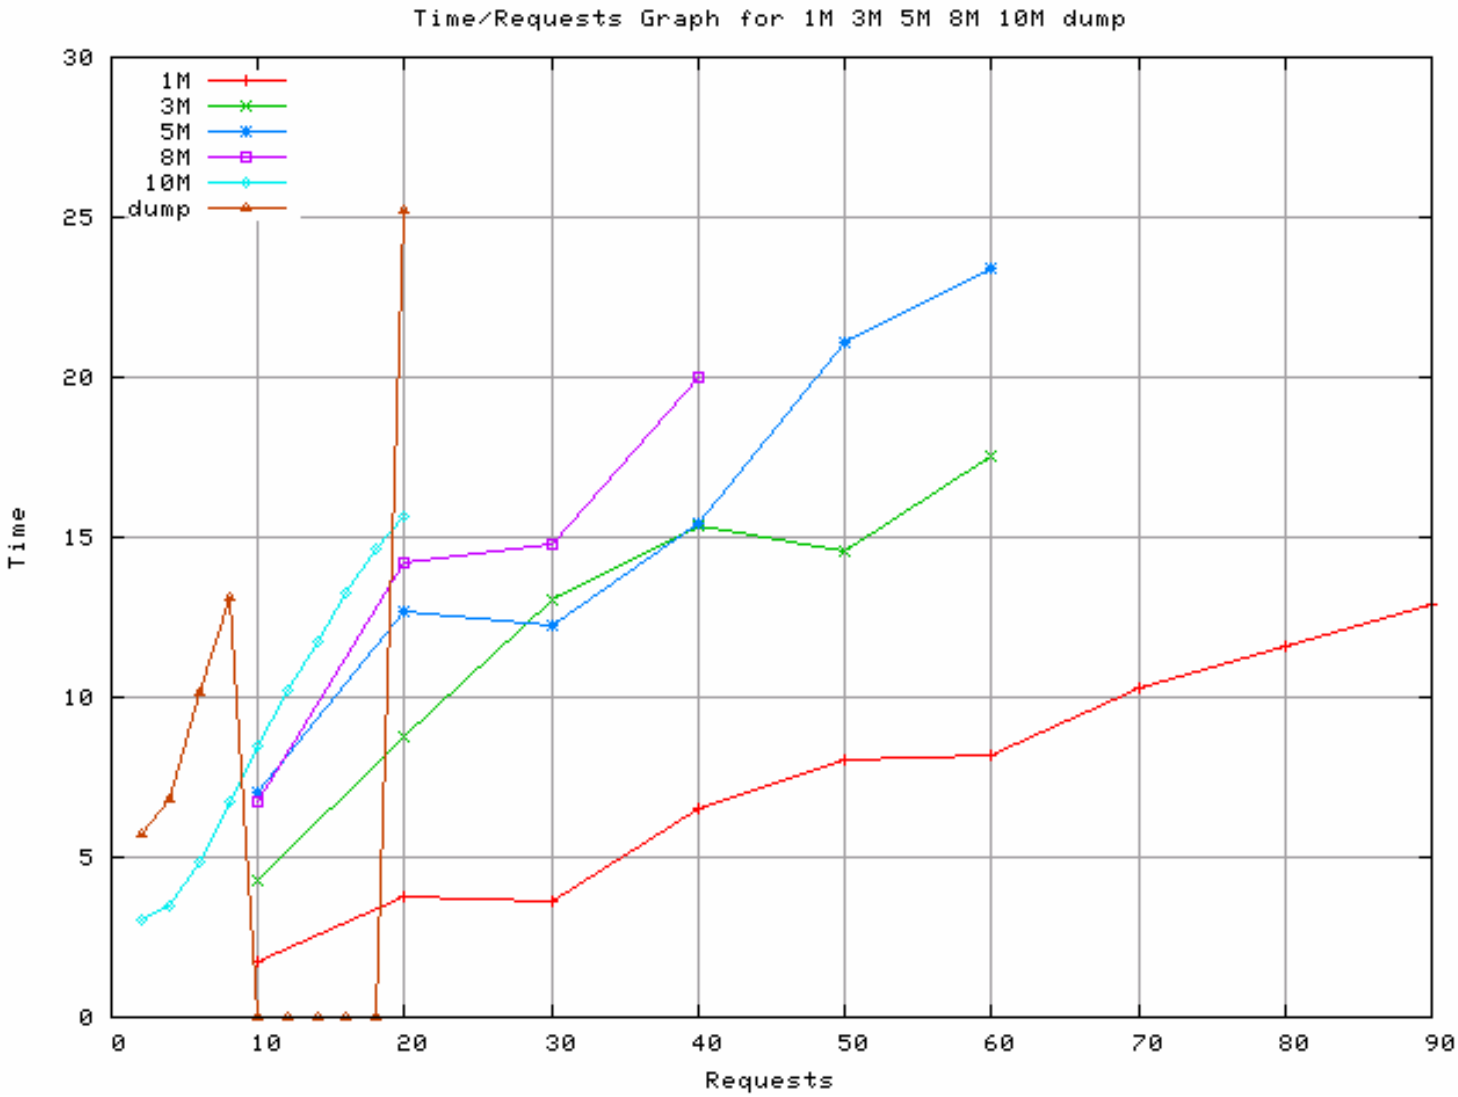

Query Scaling Of slapd

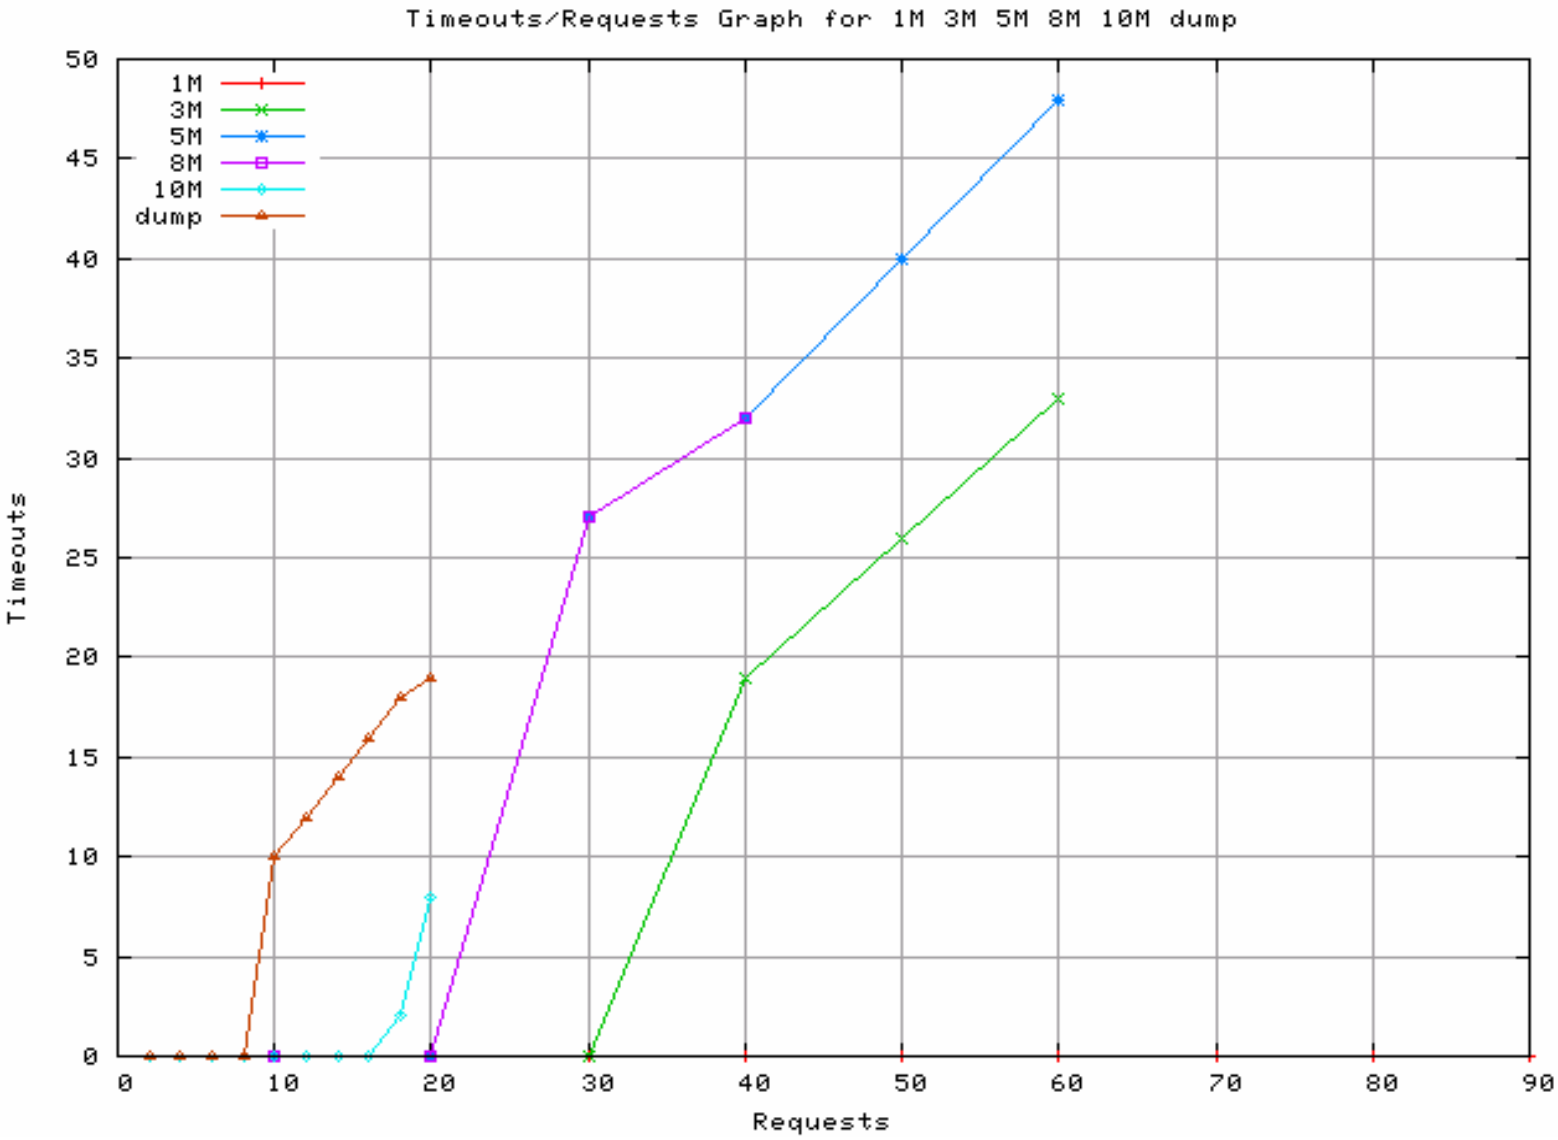

Loading Of slapd

- **Load**

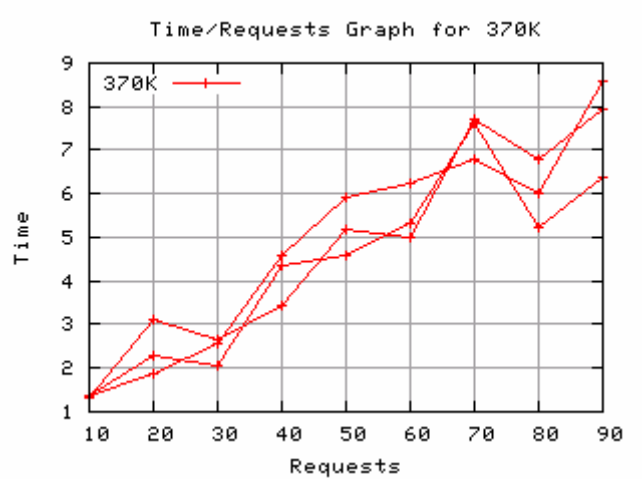

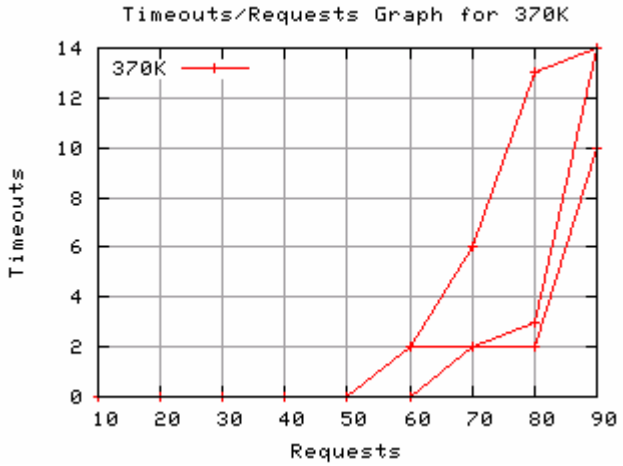

- **No load**

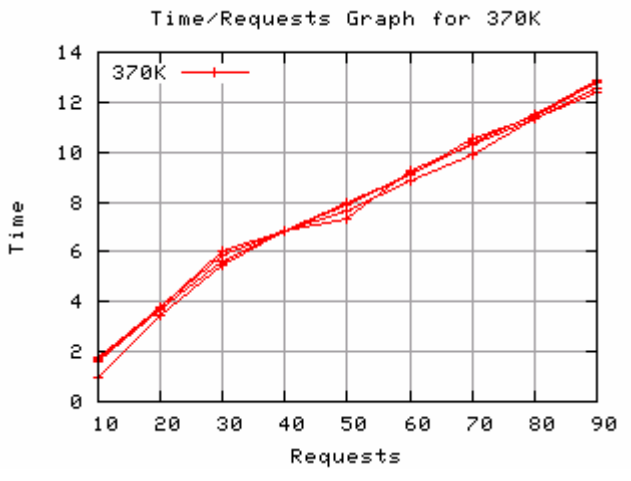

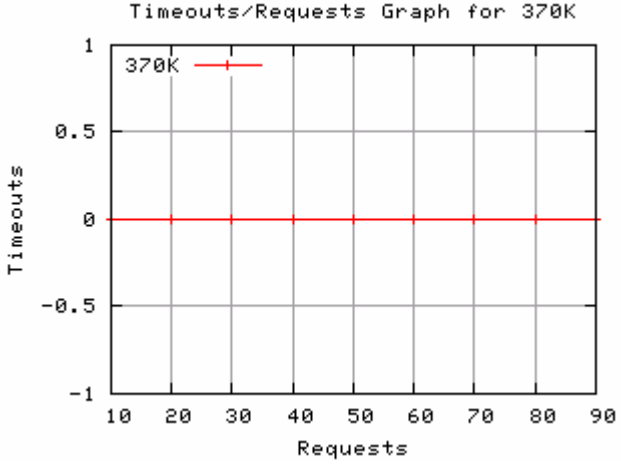

Short And Medium Term Solutions

- **Short term**
  - Put the site-level BDII on a stand alone node
  - Run the CE information provider on the site-level BDII
  - Introduce regional top-level BDIIs
    - Top spread the query load
    - Must ensure all have the same quality of service
- **Medium term**
  - Improve the efficiencies of the queries
  - Add some form of caching in the existing tools
  - Improve the query performance of the BDII service

Long Term

- **Move away from using the LDAP client directly**
  - Queries a single point and does not cache the information
- **New information client**
  - That uses caching and efficient queries
  - Use the site-level BDII for internal site queries
  - Using multiple top-level BDIIs
    - For fault tolerance and load balancing
- **Re-consider the architecture**
  - Splitting information into static and dynamic
    - Updating at different frequencies
  - More caching at the site-level
    - For common queries

Summary

- **There are two main problems**
  - High load causing the slapd not to respond
  - Time-outs on the top-level BDII
- **Investigations are still on going**
  - Scalability performance of slapd
  - Effects of high load on the slapd
- **Deployment changes will minimize the problems**
  - Run the site-level BDII on a stand alone machine
  - Run the CE information provider on the site-level BDII
  - Introduce regional top-level BDIIs
- **Clients need updating to reduce the load**
  - Improve query efficiency
  - Add caching for static information
- **Need to have a long view for scalability**
  - Re-think deployment architecture