

# A spatial DBMS buyer's guide

Maarten Vermeij<sup>(1)</sup>, Wim de Haas<sup>(2)</sup> and Wilko Quak<sup>(1)</sup>

<sup>(1)</sup> Delft University of Technology, OTB, section GIS-technology  
The Netherlands  
c.w.quak@tudelft.nl, m.j.vermeij@tudelft.nl

<sup>(2)</sup> Rijkswaterstaat, DID  
The Netherlands  
wim.de.haas@rws.nl

## Abstract

At FOSS4G 2007 we argued that benchmark results of DBMS systems is only of marginal help when choosing a spatial DBMS. Performance is only one of the few aspects. Also the high rate of development of hardware and software make the result of benchmarks only of temporary interest. In this paper we analyze what is of interest when choosing a spatial DBMS. This will include the following topics:

- The main consideration when choosing a DBMS is functionality: If the system cannot do what you want it to do it is useless. The difference between the different systems is not always clear from the documentation. We will elaborate on caveats when checking the functional requirements of a spatial DBMS against the documentation.
- Although we argued that performance is not the main issue, it is still an issue, we will hunt down when to expect performance issues and how to handle them.
- When choosing a system also hardware should be considered. Possible bottlenecks in a system are: CPU speed, memory or disk I/O. There is no use in investing in non-bottleneck components. Can the bottlenecks be predicted in advance?
- One of the choices is between open source software and proprietary software. Considerations for this choice include: Total cost of ownership (TCO) and vendor lock-in. We will discuss these issues in a spatial context.
- How well does the DBMS integrate with other components e.g. GIS systems. Also may be consider other systems that start to incorporate spatial data like ERP systems. OS software tends to support other OS software and proprietary software supports other proprietary systems. Is there a gap?
- How scalable is the product. It handles my 2 million points now, but can it handle 10,100, 1000-million as efficiently?
- Is the management open to Open Source Software or do they want to buy a 'name'?
- What is the quality of the documentation

- What are the expected future development of the product. Is there a solid base of developers. Can I expect the product still to be there in 3, 5 10 years? Or can the OS product be bought buy a competitor and taken off the market.

In the last 10 years TUDelft and Rijkswaterstaat have gained a lot of experience with many different spatial DBMS systems, this knowledge we wish to bundle in one paper. The paper can be read as a 'Spatial DBMS buyer's guide' in which all aspects of a spatial DBMS that are of importance for your choice will be handled.