

Environmental IT

Sybase's Approach to Greener Data Centres

Written by: Sarah Burnett

Published May 2008
© Butler Direct Limited

All rights reserved. This publication, or any part of it, may not be reproduced or adapted, by any method whatsoever, without prior written Butler Direct Limited consent.

► INTRODUCTION

In the past, most organisations paid little regard to the environmental aspects of their IT or the way resources were used and consumed. Now, driven by increasing awareness, costs, regulatory pressures, and equipment disposal liabilities, more companies are looking to address environmental issues in their IT strategy, along with corporate and IT governance procedures. The issue of energy consumption and heat generation, particularly in the data centre, is now officially a 'hot' topic! The cost, in terms of carbon emissions, in running IT systems has to be reduced and the IT function will have to take its own remedial action if it is not to come under increasing external legislative control.

Government legislation is playing a part in creating the problem, leading organisations to store increasing amounts of data to comply with new laws designed to promote accountability. Although these laws are, in effect, creating fuel for the data centre, it will not stop legislative bodies from demanding 'more for less'.

Reducing power consumption can be achieved through a variety of measures; more efficient software, greater utilisation or less redundancy within implemented systems, and policies that recognise that the 'green issue' is not something that is going to go away, are just the starting point. Organisations need to recognise the fact that this is not a 'fringe' issue; it has to be the central element of any implemented IT strategy.

In this White Paper, Butler Group examines Sybase's new approach to data centres and data warehouses that could help organisations reduce their data centre power consumption significantly. The approach is centred on Sybase IQ, Sybase's flagship column-based database that is designed for analytics and Business Intelligence (BI) applications.

What is the Problem?

It is a well known fact that volumes of data and information held by organisations are growing at a rapid rate. The main reasons for this are:

- The change to electronic ways of doing business, be it for transactions, services, planning and analysis, collaboration, or even today's work styles (e-Business, e-Government, mobile, and flexible working).
- More information is now being stored electronically rather than in a paper format.
- The always on and connected Web demands and generates huge amounts of content.
- Regulatory compliance requirements, such as, data protection, Markets in Financial Instruments Directive (MiFID), and Basel II require organisations to collect more information from business processes and keep it for longer periods of time than before.
- As data volumes increase so does the need for copies of data for back up and disaster recovery purposes, leading to further demand for storage.

The growth in data volumes has led to growth in demand for data centres which are heavy consumers of power. Figures from Sun Microsystems highlight the increasing power consumption of data centres – a massive 25-fold increase in six years.

Further research undertaken by Sun Microsystems, indicates that an average Data Centre burns its way through the equivalent of 80 barrels of oil or 2 tons of coal a day. At current oil prices that equates to something like US\$2.9 million a year.

Global Action Plan, a UK-based environmental group, published a report in December 2007 called 'An Inefficient Truth' that looked into the environmental impact of IT. The report showed that power consumption of IT centres between 2000 and 2006 increased by over 70% and is expected to grow by another 40% by 2020. The resulting power requirement is roughly 10% of the UK total, which Global Action Plan claims is equivalent to the output of four nuclear power stations.

However, large volumes of data are here to stay and grow. Therefore the need for handling them and realising benefits from them will not go away; making it imperative for every organisation to act and deal with the problem.

UK Governments' Response to the Problem

There are an increasing number of legislative and regulatory drivers for energy efficiency. For example, in the UK, the Government has put a climate change levy into effect; a tax on the use of energy in industry, commerce, and the public sector, with offsetting cuts in employers' National Insurance Contributions and additional support for energy-efficiency schemes and renewable sources of energy. The aim of the levy is to encourage users to improve energy efficiency and reduce emissions of greenhouse gases. Climate Change Agreements allow energy-intensive business users to receive an 80% discount from the Climate Change Levy, in return for meeting energy-efficiency or carbon-saving targets. Moreover, the UK's Climate Change and Sustainable Energy Act 2006 is designed to tackle climate change issues and secure a diverse and viable long-term energy supply. The measures in the Act are to help the UK to achieve these objectives.

Furthermore, in October 2006, the UK Government commissioned the Stern Report. Written by Sir Nicholas Stern, Head of the Government Economic Service and Adviser to the Government on the economics of climate change and development, the report examined the economics of climate change and concluded that lack of action, together with corporate and political intransigence, is likely to cost the global economy 15 times more than that needed to stabilise climate change.

These are factors that organisations must take into account, particularly as most seek to run their IT functions like a business with emphasis on service reliability and cost effectiveness.

Governments themselves face challenges in this area: they have to act in order to reduce their own carbon footprint and lead by example; they also have to balance the requirements for regulating industries and commerce which, always leads to additional data storage, against the need to stop climate change becoming an even bigger issue than it is now. Therefore Butler Group believes that Governments need to examine the effects of their legislation on the environment in their totality and as part of the legislative process.

The Exploding Data Centre and the Drivers for Change

The exponential increase in data volumes is placing enormous pressure on the enterprise in terms of costs, the infrastructure, space, and power. Energy costs have been steadily climbing over the last few years with no reversal of the trend likely for the foreseeable future. The cost of space for data centres is also going up. These are estimated to be ten times their value of just three years ago. Another imminent problem is that of the actual energy supply running out. For example, there are fears that when financial institutions in the City of London respond to MiFID in full, by deploying even more data centres, the unthinkable will happen – the City will simply run out of power!

Other drivers for change are the cost of management and administration of data storage, networks, and other infrastructure components. Together these drivers add to a compelling business case for change and confirm Butler Group's view that for too long now organisations have responded to the rise in data volumes by simply buying more hardware. Throwing more hardware at the problem is no longer a viable option and alternative approaches should be developed towards solving the problem. Butler Group believes that an accessible option is to improve the efficiency with which data is stored, accessed, and managed. The Sybase proposition aims to do just that.

► HOW SYBASE CAN HELP

Butler Group believes that the massive increase in data collection would be worthless without the ability of an organisation to use the data to deliver improved customer service and operations. In order to deliver these benefits organisations collect information into increasingly larger data warehouses or data marts, frequently breaking into multiple terabytes of storage. These data warehouses have to cope with increasing demand for analytics and their performance can be slow as a result. They also require a good deal of support and services whilst contributing greatly to the carbon footprint in the data centre.

Sybase's proposition for solving the problem is to approach it from a different angle, from the software-side; to allow enterprises to organise their data better rather than merely installing more hardware for storage. This different approach is centred on Sybase IQ, a column-based database that as the name infers, stores data values vertically rather than the row-based structures found in relational database systems. In combination the columns use sophisticated tokenisation and algorithmic techniques to remove duplication of data and compress the required storage needs.

The Design of Sybase IQ

As already mentioned Sybase IQ is a column-based database. Where as transactional systems require row-based tables, the column-based nature makes Sybase IQ particularly well suited to queries and analytics. Each column holds data of the same type and as queries often relate to large volumes of data of the same type, read processes can take place quickly. Butler Group believes that performance advantages can be achieved in complex queries for example, when large table scans are required. Sybase IQ also offers multiplexing allowing procedures to run in parallel to further accelerate query performance.

Sybase IQ tokenises the data, i.e. it uses codes to denote unique data values, and then maps those to the real data values using bit arrays. This approach changes the processing requirements from literal string comparisons that computers do not do efficiently, to Boolean operations that computers do very well. Another advantage of Sybase IQ is that by having data of the same type in each column compression it can become more efficient to gain excellent compression rates. Furthermore, the product automatically indexes data as it is loaded into the repository, which means that it does not require any indexing or specialised schemas to be applied, making the data immediately available for querying or reporting on. Sybase IQ offers a number of additional indexing options too, to speed up *ad hoc* queries.

A side effect of tokenisation and compression can be slower data load times. However, we expect the effect to be marginal and in terms of power consumption, to be more than compensated for with the lower storage requirements and the higher query performance. It is important to note that efficiency and compression rates from tokenisation may reduce if there are too many unique values to be tokenised. In the current (12.7) release Sybase IQ can tokenise columns of up to 65K values. In the 15.0 release this limit will be increased to 1.23 billion values. Conversely, efficiency and compression rates increase with tokenisation, with the natural repetition that is present in the vast majority of data. The main advantages are that query speeds increase as the volume of data reads are substantially reduced as a result of tokenised and compressed data. According to the Green Grid, a vendor consortium founded in the US with the aim of increasing energy efficiency in data centres, disks draw more power when being accessed to move the read/write head. The increase in power requirement is of the order of 30% between idle and fully utilised. Therefore more efficient read operations, as offered by Sybase IQ, can lead to a reduction in power consumption in the data centre.

Sybase claims, that in certain cases compression rates of 90% can be achieved. If achievable, that would indeed lead to significant space and power cost savings in the data centre. The Sybase case studies in the next section highlight some reductions in data storage requirements which, although not as much as 90%, are nonetheless impressive. Further information on compression rates can be obtained from Sybase, including the results of a benchmarking study carried out jointly with Sun Microsystems on a petabyte of raw data.

► CASE STUDIES

Samsung Life

Samsung Life is South Korea's well-established leader in the domestic insurance sector, where it sells individual and group life insurance through a vast branch network and more than 30,000 salespeople.

Business Challenge

The requirement was to develop a high-speed reliable process to provide field users with a single view of data. That meant integrating 10 different data marts that had previously operated independently depending on their purpose. There were three overriding goals for the project: to empower field workers to execute their own analysis; to provide access to high-conforming data through a unified view; and to reduce the amount of time for report generation.

Solution

Using Sybase IQ, the company integrated eight of the 10 different data marts and established an Enterprise Data Warehouse (EDW). Samsung Life reduced storage requirements and storage costs through Sybase IQ's data compression. Once the project goals were clear, the Samsung Life team turned to Samsung SDS, its IT expert for advice on the right solution. According to Sybase's case study, the Life Insurance Development Team at Samsung SDS concurrently reviewed the products of all database vendors including NCR, Oracle, IBM, and Sybase and ranked Sybase IQ, a database specifically designed for data warehousing, number one for its query speed and compression ratio.

Results

The group is now able to create reports rather than rely heavily on the IT function. Analytical responsibility, such as, contract status and characteristics of financial planners and customers, was switched from the IT team to the relevant business team, saving resources and improving the overall efficiency of operations.

Field workers are now able to easily access the data warehouse to extract current and accurate information in an easy-to-use environment and the company has eliminated 80% of individual data marts.

Customer service improvements were achieved due to improved speed of response and elimination of delays caused by the system.

Citigroup

Citigroup has established a new form of online historical archive in Germany, in which large volumes of data such as, years' worth of trading data, can be stored online. The system allows users to select and analyse historical data at speed, and generate reports.

Business Challenge

Legal requirements mean that listed companies in particular, have to document their business activities in more and more detail and for longer periods, as well as to store the corresponding data and to reproduce it quickly when necessary.

Solution

Citigroup has replaced its static archive with an online historical archive based on Sybase IQ.

Results

Citigroup has achieved:

- A 70% reduction in raw data storage requirement.
- The goal is to store 10 years of trading data with currently over four years of trading data stored. This amounts to 13.2 million deals and five Tera Bytes of storage.
- The system enables the organisation to analyse a growing quantity of historical data rapidly for trending and profiling purposes and also to process the data if necessary (e.g. for back testing).
- A large number of variety of analysis can be carried out on historical data, such as, statistics, trend analysis for marketing and CRM, and process optimisation.

► CONCLUSION

Butler Group believes that a more holistic approach is needed to address the effects of IT on the climate. The challenge is to square the circle of a growing data mountain and the business need to analyse data assets for business improvements. The answer may well be that we need to fundamentally change the way that we do business today, but that is highly unlikely for the foreseeable future. A more practical and realistic approach is to address environmental issues of IT in the IT strategy, along with corporate and IT governance procedures. That approach must include software power efficiency, better and more efficient data management, and more efficient query and analysis processes. Sybase's proposition fits into this area, with Sybase IQ providing a new approach that could go some way towards greener IT objectives. Sybase IQ can significantly reduce requirements for disk storage whilst keeping the data in the data warehouse available for query and analysis. Its column-based design allows more data to be read in a single read of the disk, reducing the increase in power consumption that results from disk head movements and I/O. A recent survey by Datamonitor found that demand for enterprise applications, including BI, is growing at 7% per annum. Furthermore, figures from the industry estimate that companies require increases in data storage capacity of 50%-60% per year; that is more BI requirements and more data to be analysed every year. Although the price of storage itself is coming down, the costs of power for the data centre, floor-space, cooling, and administration are not. Therefore the Sybase proposition can be good for the environment and good for corporate coffers too.

Contact Details

Sybase (UK) Limited

Sybase Court
Crown Lane
Maidenhead, Berkshire, SL6 8QZ
UK

Tel: +44 (0) 1 628 597100

www.sybase.com

Sybase

One Sybase Drive 2
Dublin 3
CA 94568
USA

Tel: +1 (0) (925)236-5000

www.sybase.com

Butler Group 
a **Datamonitor** Company

Headquarters:

Europa House,
184 Ferensway,
Hull, East Yorkshire,
HU1 3UT, UK

Tel: +44 (0)1482 586149
Fax: +44 (0)1482 323577

Australian Sales Office:

Butler Direct Pty Ltd.,
Level 46, Citigroup Building,
2 Park Street, Sydney,
NSW, 2000, Australia

Tel: + 61 (02) 8705 6960
Fax: + 61 (02) 8705 6961

End-user Sales Office (USA): Important Notice

Butler Group,
245 Fifth Avenue, 4th Floor,
New York, NY 10016,
USA

Tel: +1 212 652 5302
Fax: +1 212 202 4684

This report contains data and information up-to-date and correct to the best of our knowledge at the time of preparation. The data and information comes from a variety of sources outside our direct control, therefore Butler Direct Limited cannot give any guarantees relating to the content of this report. Ultimate responsibility for all interpretations of, and use of, data, information and commentary in this report remains with you. Butler Direct Limited will not be liable for any interpretations or decisions made by you.

For more information on Butler Group's Subscription Services please contact one of the local offices above.