

The trials and tribulations of 24/7 database management

James Cockayne, Principal DBA at Camelot Group gives us his views on key issues for DBAs, the newest active-active availability solutions and how a DBA can ensure they get their all important beauty sleep!

Just how important is it for organisations to ensure they have a robust database availability & disaster recovery solution?

The database is always the lynchpin of an application. Without the database, the application doesn't function, and there's a bunch of angry users wanting to know why (not to mention angry managers seeing money slipping down the drain). More and more there is an expectation that the services we provide will always be available when the customer wants them, and as expectations get higher the tolerance for failure gets lower. It is vital that we not only ensure the customer can access the service they want during the traditional 'working day', but increasingly all the way up to 24x7 availability.

What are some of the issues for database administrators looking after databases in a 24/7 environment?

DB2 UDB on distributed platforms offer a highly reliable database solution, but in the real world many different applications and scripts run against our databases which are hosted on various pieces of complicated hardware and problems will occur. An outage during the day is bad enough, but at least the DBAs and other support teams are at work and ready to respond. When supporting a 24x7 system failure in the middle of the night can be, literally, a nightmare as the DBA has to get connected and conduct a diagnosis and fix as quickly as possible – after all, in a 24x7 system the middle of the night here in the UK is daytime for customers elsewhere in the world. Often an issue will mean calling (and usually waking up) other support personnel from teams such as server support, storage, or networks for diagnosis or resolution - not to mention managers on the escalation list...

So how can those midnight calls best be avoided?

For both happier customers and support staff a backup standby server with an automated failover solution is often the answer. This isn't a panacea for all database woes - but in a well run and maintained production DB2 environment a sizable percentage of any unplanned downtime will be caused by hardware or third party software issues, so having the database automatically move onto another server can make a serious difference to the time it takes to get the applications up and running again.

So, will end users still experience an outage with an automated failover system?

Depending on the exact solution, and the application the database supports, the time to get the production service available again could be anything from a matter of minutes if using features such as DB2's HADR in conjunction with Automatic Client Reroute, up to an hour if relying solely on mechanisms such as HACMP and scripting to start up DB2 on a cold server and restart the applications. Compare this to potentially several hours of work to manually restore a database onto a standby server and reconfigure so the application can see it, and the benefit is clear – it's always worth setting the management expectations in advance though: it's automatic, not instant!

You mention HADR and HACMP as potential solutions, what are the main differences?

HADR is a warm standby solution and gives not only fast failover but also high availability while applying DB2 FixPaks, effectively meaning the only downtime to install maintenance is switching between primary and standby servers. However it needs twice as much disk space because your database requires both the primary and standby servers have their own copy of the database; a pure HACMP failover solution will use a cold standby and so can share one copy of the database on disk.

So, the database automatically fails-over to the standby server. Great! That means you can carry on your beauty sleep undisturbed, right?

Not necessarily, automatic failover isn't the answer to all database problems – indeed if your database failover kicks into action the DBA and other support teams will still need to diagnose and fix the problem, and unfortunately it's still the case that things will go wrong which can't be solved by failover (batch jobs will still error and require sorting out, user error will certainly still occur...) so this isn't an end to phone calls in the middle of the night. What automatic failover gives us is a faster response to getting the production systems available to the users again, and enables fault diagnosis and resolution to take place possibly at a more sociable hour - certainly at a more careful and considered pace.

What benefits would you see from an active-active high-availability solution?

Having an active-active solution can provide the ultimate in high availability for DB2 UDB on distributed platforms. Whilst traditional thinking is based around minimising the time taken for the database to become available on a standby server after a serious problem on the primary, with active-active the standby server is already active – using a virtualisation product it's possible for a failure to be invisible both to the users and even the application itself. Couple this with the ability to perform maintenance on one server whilst the other(s) continue to be available to the application and it's realistic to be aiming for 100% availability.

There are also other benefits to going for an active-active solution which are especially relevant at this time: in failover solutions there is a standby server which has purchase and maintenance costs but is effectively only used in an emergency. With active-active this server can be providing more of a return on that investment - by taking on some of the workload which would otherwise all be routed through the one 'primary' server performance can be improved. In a system with a growing workload utilising this solution could even save the cost of upgrading the hardware to cope with demand.

About Triton

Triton Consulting are Database Management Specialists and IBM Premier Business Partners providing cost effective solutions to organisations of all sizes. With our Consultancy on Demand offering you can purchase between 10, 20 or 40 hours of consultancy which can then be used as and when you need it. Hours can be used for training, development work or to provide cover during sickness and holiday within an existing DBA team.

For more information visit www.triton.co.uk