



## Re: Data Protector 6.0 scheduling mysteries

weekends, etc.

Now that some incremental backups just take 4 minutes or so, the created schedules would be quite tight. For example:

17:00 A  
17:00 B  
17:10 A  
17:10 B  
17:30 C  
17:30 D  
17:40 B  
17:40 D

A to D are backup specifications and the first occurrences are incremental backups, followed by levelled backups.

As DP cannot schedule at 17:10, the backup was scheduled at 17:00 also.

The expectation was that it would be queued until the first one is finished.

However, as it seems, DP just silently ignored the second occurrence of

A,  
B,  
C, and D.

Using the standard DP interface, can you actually schedule the one specification to run twice at the same time? I have not tried, but assume the GUI would "merge" both of them into one.

They (17:00 and 17:10) are not actually the same: They differ in backup level and data protection.

Yes, but they use the same "specification", so DP seems to think they are the same. Just like DP does not understand that the same data backed up via different specifications are just different versions, not completely different data.

Re: Data Protector 6.0 scheduling mysteries

Can anybody explain this arbitrary restriction to multiples of 15 minutes?  
Cron can do a better job for 20 years now.

The omnitrig process is what actually schedules the sessions, and it is only run every 15 minutes (from cron). I suspect this has been hardcoded in somewhere.

I guess they have a fixed-size array to manage the schedule somewhere, and they wanted to keep the size of the array small. Wrong design.

You want a list of "wrong design" choices with DP ... I don't have enough bandwidth to post that. ;)

You can however run a session "manually" from the command line ... or from cron. So you can just ignore DataProtector's scheduling stupidity, and schedule them directly from cron. Not as pretty, but if you have already written a scheduler, formatting it's output for cron (vs DP) should not be too difficult. Look at "omnib -datalist specname -no\_monitor"

Actually I'd be using at(1) then, but the next surprise could be what happens when I schedule about 3000 jobs (for one year) using at(1).

The solution seems to be to specify a minimum duration of 15 minutes per backup session, even though the backup just needs a few minutes. This will fragment the time space a lot, being unable to fill the gaps. Also this will minimize device usage as there are unnecessary breaks

Re: Data Protector 6.0 scheduling mysteries

between backups.

My scheduler could handle free time slots down to one second, but it does not make any sense with DP.

Be aware that you will probably hit another quirk with DP queuing. When you have multiple sessions queued (for a device or licence) DP does not respect queue order. Once the device (or licence) becomes available there does not seem to be any rational (except for Murphy's Law) as to which session will proceed first.

Maybe someone should tell those programmers what semaphores are used for.

There are dozens of ways for them to fix this problem, but they just don't see it as an issue. I often get sessions timeout because a bunch of "later" sessions keep getting in front of it on the "queue".

BTW: How do you guys keep track when which backups are scheduled to avoid queueing conflicts caused by scheduling and device usage?

Thankfully (for me) I have a small number of long (full backup) sessions, so scheduling is not too difficult.

Yes, a daily full backup of everything would also solve the problem, but it's a waste of material.

In my environment it's actually the only way I can even get a backup. I'm

Re: Data Protector 6.0 scheduling mysteries

not complaining too loudly ;)

--

Posted via a free Usenet account from <http://www.teranews.com>

.