

Is Parallel Repair Really Better Than Sequential Repair?

April 2008

We recently had a thought-provoking challenge from one of our readers questioning our reasoning on parallel repair versus sequential repair.¹ In previous articles in the *Availability Digest*, we had argued that parallel repair would restore a downed system faster than sequential repair.

Sequential and Parallel Repair

Under sequential repair, there is only one service technician. In a multinode system with one spare, if a node fails, the technician begins repairing that node. Should a second node fail while the first node is still under repair, the system fails. The service technician continues his repair of the first node; and when he has returned that node to service, the system is restored to service. He then begins the repair of the second failed node.

Under parallel repair, there are two service technicians available. Should a first node fail, one of the service technicians will start to repair that node. If a second node fails while the first node is still being repaired, the system fails. However, as opposed to sequential repair, as soon as the second node fails, a second service technician begins the repair of that node. The system is returned to service as soon as either one of the failed nodes is repaired.

The Parallel Repair Advantage

We argued that the rate of repair of failed nodes under parallel repair would be twice as fast as that under sequential repair because there were two service technicians performing repairs. If the repair rate were twice as fast, the average time to repair the first node would take half as long. Therefore, a failed system would be returned to service in half the time under parallel repair as compared to its restore time under sequential repair. If mtr is the mean time to repair a single node, and if R is the average system restore time, then

$$\begin{aligned} R &= mtr && \text{for sequential repair} \\ R &= mtr/2 && \text{for parallel repair} \end{aligned}$$

We extended this argument to a system with s spares and $(s+1)$ repair technicians. In this case, parallel repairs would occur at a rate that is $(s+1)$ times faster than that of a single repair technician; and the average system restore time would be

$$R = mtr/(s+1) \quad \text{for } s \text{ spare nodes and full parallel repair}$$

¹ [Calculating Availability – Repair Strategies](#), *Availability Digest*, November, 2006.
[Failure State Diagrams – Repair Strategies](#), *Availability Digest*, October, 2007.

A Question for the Editor

In response to this reasoning, Alan from New Zealand wrote:

"Hi Bill, The Availability Digest arrived at a most interesting time - the hard disk on my machine has just died (literally) - and it raises interesting aspects of my own approach to backup and availability!

However, I have a question for you about your rules of repair. You state that if it takes a repair crew four hours to repair a node, then they can repair two nodes per eight-hour day. Two crews can repair four nodes per day, so the average time to repair a node becomes two hours. Now project management tells me this is not true. What is true is that the average time of repair becomes two hours, but the elapsed time for the first repair remains at four hours. So this adds a level of complexity to availability, does it not? I am unsure of the math that follows, but I do know in project management terms that if I assume that the time to effect a repair is four hours, and I reduce the entire elapsed time for repairing N nodes to four hours total by adding N repair crews, my average time per repair becomes $4/N$. However, if all of the crews started at the same time, the first node still becomes available four hours later, but I would get N nodes available at that time. If I add $N/2$ crews, I get the repair done in eight hours for all nodes. The average repair rate is $8/N/2$, but the first node to be repaired is still four hours later.

This issue is the one that many people in project management make a mistake with - they confuse resources with planned time to effect a task. For example if I have to dig ten holes, in the plan I should allow the time for only one hole to be dug, but I need ten of them. The resource allocation controls the time for the project to be completed, and I can shorten the time for the project to be completed by adding more resources for the tasks (in this case adding more hole-digging crews).

So my question is how does this elapsed time affect the math in your availability measure? It seems to me that availability is not improved, but yes, MTR average is improved. Yes, my availability is improved by increasing the number of repair crews in that I get more failed nodes available again within the elapsed time, but the actual time for repair remains four hours in the example you gave. In a chaotic situation where nodes are failing randomly (the real world, I guess), the average may apply as only one crew can work on a failed node so they will start at the time of failure. The second node fails some time later, and the second crew starts work, etc. On the basis that when nodes come back on line at about the same time that another fails, there may only be, say, two nodes out at any one time and it would appear that the MTR is two hours, but the node repair time does stay at four hours. But there is some effect of rate of failure and randomness of those failures on MTR as I see it."

My response to Alan follows:

"Alan - I can see that you've really thought about this, and the point that you raise is quite interesting. There is an answer to your quandary, and it is in the mathematics and has to do with the probability distribution of repair time.

If repair time is constant - always four hours in your case - then you are quite right. It will be four hours until the first system is repaired. However, repair times aren't generally constant. Some can be accomplished quite rapidly, and some might take some time. If the failure is correctable by simply replacing a customer-replaceable component that you happen to have on site, repair might take just a few minutes (plus reboot time). To have to order a part and install it when it arrives might take a day or more. So if two nodes fail, it is likely that one will be a faster repair than the other; and that is the one that will be fixed first if parallel repair is used.

So you have to assume some sort of probability distribution of repair times. The case that I discuss (though I didn't say it and probably should have) is the simplest case mathematically (short of a constant repair time), and that is that repair times are governed by an exponential distribution.

Whether that is the case or not, parallel repair will be faster than sequential repair if repair times are not constant. With sequential repair, the first node that failed is worked on first. If it is the one that takes the longest time, too bad. That will be the repair time.

If parallel repair is used, then whichever node is repaired the fastest will be the one to come up first. Though the precise numbers may change, I think the conclusion that parallel repair is faster than sequential repair is a valid observation.

Perhaps a better model for sequential repair would be one that, upon the failure of the second node, repair stops on the first node; and the second failure is analyzed. Then whichever repair is deemed to be the shortest is undertaken. However, in the heat of a crisis, this might not be feasible. Also, often times the analysis of the fault is what takes the most time. The actual repair might be quite fast compared to the time that it takes to decide what had happened.

In any event, I'm glad that you read the Digest and that it stirred some discussion and thinking. Let me know what you think about the above argument. "

Alan concludes this conversation with the following:

"Thanks, Bill - you are quite right about the probability distribution and the analysis time. I like the reply very much; and it is, of course, correct. The fault on my part was to read the article literally and not think about the actual failures, time to repair, ability to replace modules, the parallelism that applies with multiple repair crews, etc., etc. (the normal chaotic systems we work in all the time - chaotic in this sense = unpredictable).

I do read (and enjoy) the Digest. It adds a higher dimension to the availability issues I deal with in the small/medium (by USA standards very, very small) companies and situations in NZ. For high availability solutions, I am usually dealing with paired machines that are not always in lock step for licensing reasons (sometimes people just won't spend money where they should) so people accept an RPO of two to five minutes with an RTO of up to six hours. When it happens, however, they want the RTO to be MUCH shorter - next purchase round we usually have a much simpler time of handling the budget! I also have to say that their transaction rate is such that an RPO of two to five minutes is realistic - coupled with their ability to repeat any lost transaction."

Thanks, Alan.

So What Is the Real Story

Alan's comments got me thinking. What really is the relationship between the distribution of repair times and the average system restore time under parallel repair? So, with a little math (which is detailed in the document [Analysis of Repair Strategies](http://www.availabilitydigest.com/public_articles/0304/repair_strategy_analysis.pdf), available at www.availabilitydigest.com/public_articles/0304/repair_strategy_analysis.pdf²), I was able to make the following observation:

As the distribution of repair times morphs from a constant distribution to an exponential distribution and beyond, the advantage of parallel repair over sequential repair increases.

² which I recommend only for hard-core math nuts,

This observation is illustrated below. Let us define a repair advantage A :

$$A = R_s/R_p$$

where

R_s is the mean time to restore the system under sequential repair.

R_p is the mean time to restore the system under parallel repair.

That is, A is the ratio of improvement in system restore time of parallel repair over sequential repair. As seen in the diagrams below,

- $A = 1$ for constant repair times (there is no advantage of parallel repair over sequential repair – both lead to the same average system restore time).
- $A = 1.5$ for a uniform distribution of repair times (that is, it is equally likely that the repair will be finished at any time up to a maximum repair time).
- $A = 2$ for an exponential repair time distribution.
- By inference, A increases further as the repair time becomes closer to a constant repair time of zero (this is not verified by calculation).

