

Podcast Episode - The Long Awaited Index Episode

For a couple of months we have had requests on the SQL Data Partners Podcast to do an episode on indexing. It is my pleasure to announce that the Indexing episode ([episode 81](#)) was released today, well at least half of it. After we finished the recording, we realized that there was too much content to fit it in to our normal episode length, so we split it into a Part 1 and Part 2 episode. Part 1 launched this week, and Part 2 will come out next week.

The guest this week was Randolph West (Twitter: [@BornSql](#)) a SQL Consultant at BornSQL, where we does performance tuning optimization, works with best practices for DR and SQL Server maintenance. Randolph has also recently become a Microsoft MVP. Between Randolph, Carlos and I we were able to cover an extensive amount of detail on indexing.

In Part 1 we covered:

Using the phone book as an example, we chat with our guest Randolph West and explain heaps, non-clustered indexes, and clustered indexes. We also explain best practices to use them when creating new tables. We also cover the issues with using wide data types in your clustered indexes, along with the drawbacks of GUIDs (Uniquelidentifiers) as a clustered index.

We didn't forget to cover the problems you might run into. For example, you might not think about indexes until your data grows very large. Sure, you can attempt to resolve the problem by moving to SSD's and increasing memory, but this will push the issue further in the future. In this episode, we will discuss considerations for indexes and why database architecture is so important. We would like to hear your opinions. What are your strategies when using indexes? Use the hashtag #sqlpodcast and let us know!

We had a lot of fun with this episode. I hope you enjoy it.

Episode Quote

“The magic word you use there was fragmentation. This used to be a huge problem as Steve mentioned with spinning drives. While I disagree with the fact that I’m not worrying about fill factor at all is ok. I disagree, you should worry about fragmentation. You should worry about these things purely from the point of views of knowing that your database isn’t in a good condition is well maintained.”

Listen to Learn

- Heaps and indexes types explained
- Randolph’s thoughts on clustered indexes vs. non-clustered indexes
- Similarity between clustered indexes and primary keys
- The results of using Management Studio Designer for creating new tables
- How to create covering indexes
- The negatives of having too many indexes and what you can do about it
- How are default fill factors affecting indexes? Is faster storage good reason to avoid fixing the indexes?

Related Links

- Randolph on Twitter [@BornSql](#)
- [Episode 81 show notes](#)
- [SQL Data Partners Podcast](#)