Attendee	Questions Asked by Attendee	Answers by Mike Cain
Alan	Did I see you using iSeries Navigator to get to Index Advice?	Yes, the examples are from System i Navigator, Run SQL Scripts and Visual Explain
Alan	Is there a way to use Index Advice on an IFS system?	No, index advice is based on relational database tables and queries
Chuck	Sometimes indexes are not recommended but we can clearly see that an index is needed. once we create it, the optimizer uses the newly created index. Why?	There are limitations to what the optimizer will advise for indexes, and when the indexes are advised. In addition to the DB2 tools, it is best to use your own knowledge and experience to determine the indexing strategy.
Chuck	Is having too many indexes good or bad?	Neither. For queries (SQL statements with local selection, joins, grouping) you must first determine and implement a good indexing strategy, then worry about index maintenance. Avoiding full table scans, temporary indexes as well as giving the query optimizer statistics are key to efficient processing and good performance.
Chuck	Can I create a derived index to replace mat's What new improvement in indexing are in the future release?	In 6.1 you can create a derived index with SQL. Example: CREATE INDEX my_schema.my_index ON my_schema.my_table (UPPER(col1)). See the SQL reference guide for more information.
Esben	Is it possible to estimate the size of a recommended index?	Technically yes. The equation is complex with many variables. Factors include the size of the key, number of keys and various attributes of the index. One technique is to evaluate your own indexes comparing the index / keyed LF size to the table / PF size and noting the number of key columns / fields. Remember that select/omit and a WHERE clause on the index will potentially reduce the number of keys in the index.
Genyphyr	Hello, one of the questions I often get from clients is 'how do I know if I have created 'too many' indexes' and the system is taking longer to maintain them than the value they bring - what are some of the performance tools we can use and what might we look for in these tools to avoid creating 'too many' indexes on our database ? I am thinking of high transaction ERP systems for example. Also has the rebuild time for EVIs when the index grows past it's originally allocated size been improved or changed at V6R1 and again at V7R1 and where is this information documented so I can do more research on EVIs to help me know when it is appropriate to build one or not.	You must first determine and implement a good indexing strategy, then worry about index maintenance. Avoiding full table scans, temporary indexes as well as giving the query optimizer statistics are key to efficient processing and good performance. Determining if INSERT, UPDATE and DELETE operations are being hampered by index maintenance requires more detailed analysis of the performance and is out of the scope of this session. Buiding EVIs is based on understanding the data, the queries and the query plans. By following the general guidelines for the use of EVIs, index maintenance will not be an issue. It is best to get educated on DB2 for i SQL query performance monitor, analysis and tuning.

Attendee	Questions Asked by Attendee	Answers by Mike Cain
Glynn	Are there any guidelines on index creation and the amount of system resources required to maintain those indexes	No, only that you must first determine and implement a good indexing strategy, then worry about index maintenance. Avoiding full table scans, temporary indexes as well as giving the query optimizer statistics are key to efficient processing and good performance. Determining if INSERT, UPDATE and DELETE operations are being hampered by index maintenance requires more detailed analysis of the performance and is out of the scope of this session.
Glynn	Is index maintenance quicker when there are no deleted records in the physical file? Question should have been - will logical files build quicker if the physical does not have deleted records?	Regarding maintenance: no, I do not believe so. Deleted rows take up space and must be read during the index / keyed LF build process. Deleted rows are not processed. Have no or few deleted rows will allow the full table scan to perform faster.
Henry	Sometimes the index analyzer recommends an index that already exists. Why?	This should not happen unless the index being recommended has different attributes than the existing one.
Keith	Is there a method to segregate indexes advised for a particular batch run, end of day processing, etc?	I believe you could use some clever monitoring and querying of the index advice table to gather the information. Otherwise you can make use of the SQE plan cache and/or SQL performance monitor to capture information by time period, job, user, etc. These tools also capture index advice information from the optimizer.
Lars	What will solid state disks do for performance?	SSDs can help to improve random read operations. Indexes, keyed logical files are candidates for SSD storage, as well as some table and physical files that are the targets of random look ups. Proper monitoring, analysis and planning are keys to determining where and when SSDs
Lindsay	Now V6.1 .... still large tables .... still running with REOPTIMIZE_ACCESS_PLANS *ONLY_REQUIRED .... how can I turn the "advisor" off and on.	:-) You cannot turn the advisor off as it is part of the query optimization process. Periodically review and then delete the index advice information located in the table.
Mike	For a database with RI at the third normal form, should EVIs be created for the Foreign Keys of the Child Tables?	It depends on the actual data model, data and queries. For star and snowflake schema models, this is a common strategy to support DB2 for i "look-ahead predicate generation" or LPG. http://www.mcpressonline.com/database/db2/the-power-and-magic-of-lpg.html
Mike	Are EVI better for batch processing verses on-line processing.	It depends on the actual data model, data and queries. EVIs have certain advantages when it comes to indexing ANDing/Oring and index only access. A good understanding of query optimization and DB processing techniques are important when considering the creation of EVIs.

Attendee	Questions Asked by Attendee	Answers by Mike Cain
Naveen	While creating a logical files , We have option to Select or Omit Records. Is this achievable with Indexes? Is SQL View the only available alternative?	Yes, starting with 6.1 the SQL CREATE INDEX and CREATE ENCODED VECTOR INDEX statements allow a WHERE clause to be specified. SQL views contain no data, thus they are not used to access the table - in other words, the view is really just a query definition.
Rahul	Hi Mike, my question is it possible to share any bench mark figures on workload increase on CPU utilization before and after indexing?	The difference is huge and can be demonstrated very easily in your own shop. Build a simple data model of 3-5 table that are related. Run some queries with local selection and joins against the tables. Then create the advised indexes and run again. Measure response time as well as resource utilization.
Reid	Proactively, would it be a good idea to create indexes on select and where clauses?	Yes, local selection and join columns are the primary candidates for indexing.
Roger	Is cardinality high or low relative to the total number of rows, or an absolute metric?	High or low cardinality is relative and it does take into consideration the number of rows or number of duplicates. 1,000,000 unique rows represents high cardinality for the column that is unique. 1,000,000 duplicate rows represents low cardinality. This concepts helps us understand the data and understand the choices made by the optimizer.
Roger	Often we get recommendations to add a particular index (when doing a Visual Explain, for example), but adding the index makes no difference to the query, which ends up not using it. So why was it recommended in the first place?	I covered this in the presentation. With no indexes, the optimizer is limited in its understanding of the data and the selectivity of the query. Adding the index provides the optimizer statistics and thus a better understanding of the data. Example: 1,000,000 rows and no indexes. SELECT * FROM T1 WHERE C1 = 123. The optimizer assumes C1 = 123 is selective and will produce few results. The optimizer chooses a full table scan as no indexes exist. It recommends an index on C1. After creating the index, the optimizer now understands that all 1,000,000 rows contain C1 = 123 and a full table scan is still the best choice. The index was used for statistic and not used for implementation.
Roger	For very small tables - say 100 rows or so, or fewer, are indexes still useful, or will the system just hold the table in memory anyway.	If these small tables are used in join queries, it will be advantageous to have indexes in place.

Attendee	Questions Asked by Attendee	Answers by Mike Cain
Roman	If a plan does not exist and the system picks a poor performing plan, can the system pick or change the plan based on statics information that is gathered over time?	Yes, the query optimizer will look for changes in the system resources, query environment and data. In 7.1 "adaptive query processing" is available. AQP allows the optimizer to watch the query while it's running. If not results are produced and the query is not making progress, the optimizer can look for a better plan and replace the currently running plan. All without intervention by the user or awareness of the application.