100% Pass 2025 Oracle 1z1-084: High Hit-Rate Oracle Database 19c Performance and Tuning Management Exam Outline

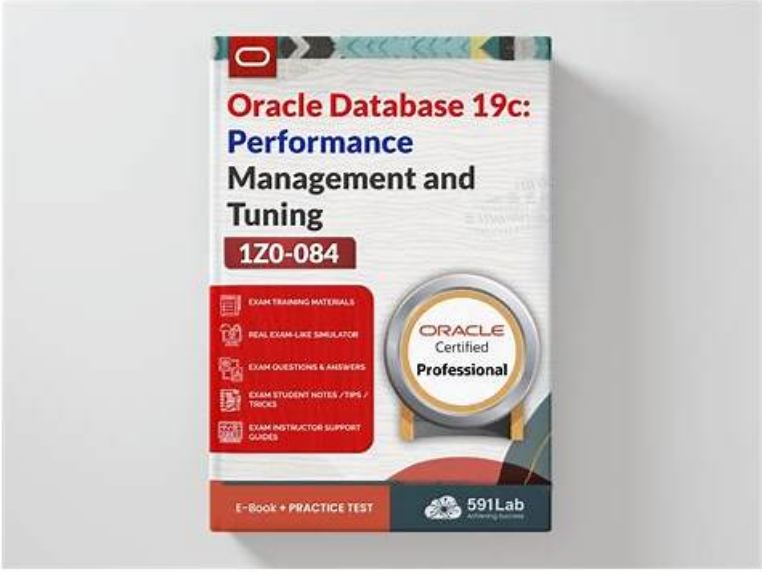

BONUS!!! Download part of Real4exams 1z1-084 dumps for free: https://drive.google.com/open?id=12zmuODnh_5V_atFMNTiAKgC5-rqy-Zr2

By browsing this website, all there versions of 1z1-084 practice materials can be chosen according to your taste or preference. In addition, we provide free updates to users for one year long. If the user finds anything unclear in the 1z1-084 practice materials exam, we will send email to fix it, and our team will answer all of your questions related to the 1z1-084 practice materials. If the user fails in the 1z1-084 practice exam for any reason, we will refund the money after this process. We promise that you can get through the challenge within a week.

Oracle 1Z0-084 is a certification exam that tests the performance and tuning management skills of individuals working with Oracle Database 19c. 1z1-084 exam is designed to assess the knowledge and understanding of the concepts and practices used in managing and optimizing the performance of Oracle databases. Candidates who Pass 1z1-084 Exam earn the Oracle Database 19c Performance and Tuning Management Certified Expert credential.

>> 1z1-084 Exam Outline <<

Latest 1z1-084 Exam Price & 1z1-084 Official Practice Test

Different with other similar education platforms on the internet, the Oracle Database 19c Performance and Tuning Management guide torrent has a high hit rate, in the past, according to data from the students' learning to use the 1z1-084 test torrent, 99% of these students can pass the qualification test and acquire the qualification of their yearning, this powerfully shows that the information provided by the 1z1-084 study tool suit every key points perfectly, targeted training students a series of patterns and problem solving related routines, and let students answer up to similar topic. It may say, the 1z1-084 Test Torrent can let users in a short time, accurately grasp the proposition trend of each year, doing all effects in the process of the difficulties in the hot, user's weak link and targeted training, and exercise the user's solving problem ability, eventually achieve the objectives of the pass Oracle Database 19c Performance and Tuning Management qualification test.

Oracle Database 19c Performance and Tuning Management Sample Questions (Q34-Q39):

NEW QUESTION # 34
Database performance has degraded recently.
index range scan operations on index ix_sales_time_id are slower due to an increase in buffer gets on sales table blocks.

Examine these attributes displayed by querying DBA_TABLES:

```
OWNER                          TABLE_NAME                     NUM_ROWS     BLOCKS EMPTY_BLOCKS  AVG_SPACE
 CHAIN_CNT AVG_ROW_LEN
------------------------------ ------------------------------ ---------- ---------- ------------ ----------
---------- -----------
SH2                            SALES                            3675372      18337            0          0
         25          37
```

Now, examine these attributes displayed by querying DBA_INDEXES:

```
OWNER                          TABLE_NAME     INDEX_NAME                     BLEVEL LEAF_BLOCKS DISTINCT_KEYS
AVG_LEAF_BLOCKS_PER_KEY CLUSTERING_FACTOR STATUS    NUM_ROWS
------------------------------ ------------ ------------------------------ ------ ----------- ------------- --------
---------------------- ----------------- -------- --------
SH2                            SALES          IX_SALES_TIME_ID                    3      353967       3541460
                     9           2830691 VALID    3675372
```

Which action will reduce the excessive buffer gets?

- A. Re-create the SALES table sorted in order of index IX_SALES_TIME_ID.
- B. Re-create the SALES table using the columns in IX_SALES_TIME_ID as the hash partitioning key.
- C. Partition index IX_SALES_TIME_ID using hash partitioning.
- D. Re-create index IX_SALES_TIME_ID using ADVANCED COMPRESSION.

Answer: D

Explanation:
Given that index range scan operations onIX_SALES_TIME_IDare slower due to an increase in buffer gets, the aim is to improve the efficiency of the index access. In this scenario:
* B (Correct):Re-creating the index usingADVANCED COMPRESSIONcan reduce the size of the index, which can lead to fewer physical reads (reduced I/O) and buffer gets when the index is accessed, as more of the index can fit into memory.
The other options would not be appropriate because:
* A (Incorrect):Re-creating theSALEStable sorted in order of the index might not address the issue of excessive buffer gets. Sorting the table would not improve the efficiency of the index itself.
* C (Incorrect):Using the columns inIX_SALES_TIME_IDas a hash partitioning key for theSALES table is more relevant to data distribution and does not necessarily improve index scan performance.
* D (Incorrect):Hash partitioning the index is generally used to improve the scan performance in a parallel query environment, but it may not reduce the number of buffer gets in a single-threaded query environment.
References:
* Oracle Database SQL Tuning Guide:Managing Indexes
* Oracle Database SQL Tuning Guide:Index Compression

NEW QUESTION # 35
You must write a statement that returns the ten most recent sales. Examine this statement:

```
SELECT s.prod_id, p.prod_name, s.amount_sold, c.cust_name, c.cust_city
FROM   sales s, customers c, products p
WHERE  s.cust_id = c.cust_id AND
       s.prod_id = p.prod_id
ORDER BY s.time_id
FETCH FIRST 10 ROWS ONLY;
```

Users complain that the query executes too slowly. Examine the statement's current execution plan:

```
| Id | Operation              | Name      | Starts | E-Rows | A-Rows | A-Time      | Buffers | Reads | OMem | 1Mem | Used-Mem |
|  0 | SELECT STATEMENT       |           |      1 |        |     10 |00:00:00.32 |    3130 |     3 |      |      |          |
|* 1 | COUNT STOPKEY          |           |      1 |        |     10 |00:00:00.32 |    3130 |     3 |      |      |          |
|  2 | VIEW                   |           |      1 |   918K|     10 |00:00:00.32 |    3130 |     3 |      |      |          |
|* 3 | SORT ORDER BY STOPKEY  |           |      1 |   918K|     10 |00:00:00.32 |    3130 |     3 | 2048 | 2048 |  2048 (0)|
|* 4 | HASH JOIN              |           |      1 |   918K|   918K|00:00:00.22 |    3130 |     3 | 1250K| 1250K| 1579K (0)|
|  5 | TABLE ACCESS FULL      | PRODUCTS  |      1 |    72 |     72 |00:00:00.01 |       3 |     3 |      |      |          |
|* 6 | HASH JOIN              |           |      1 |   918K|   918K|00:00:00.15 |    3126 |     0 | 4696K| 1834K| 4597K (0)|
|  7 | TABLE ACCESS FULL      | CUSTOMERS |      1 |  55500 |  55500 |00:00:00.01 |    1521 |     0 |      |      |          |
|  8 | PARTITION RANGE ALL    |           |      1 |   918K|   918K|00:00:00.10 |    1604 |     0 |      |      |          |
|  9 | TABLE ACCESS FULL      | SALES     |     28 |   918K|   918K|00:00:00.10 |    1604 |     0 |      |      |          |

Predicate Information (identified by operation id):

1 - filter(ROWNUM<11)
3 - filter(ROWNUM<11)
4 - access("S"."PROD_ID"="P"."PROD_ID")
6 - access("S"."CUST_ID"="C"."CUST_ID")
```

What must you do to reduce the execution time and why?

- A. Create an index on SALES.CUST_ID to force an INDEX RANGE SCAN on this index followed by a NESTED LOOP join between CUSTOMERS and SALES.
- B. Replace the FETCH FIRST clause with ROWNUM to enable the use of an index on SALES.
- C. Create an index on SALES.TIME_ID to force the return of rows in the order specified by the ORDER BY clause.
- D. Collect a new set of statistics on PRODUCT, CUSTOMERS, and SALES because the current stats are inaccurate.
- E. Enable Adaptive Plans so that Oracle can change the Join method as well as the Join order for this query.

Answer: C

Explanation:
The execution plan shows a full table access for the SALES table. To reduce the execution time, creating an index on SALES.TIME_ID would be beneficial as it would allow the database to quickly sort and retrieve the most recent sales without the need to perform a full table scan, which is I/O intensive and slower. By indexing TIME_ID, which is used in the ORDER BY clause, the optimizer can take advantage of the index to efficiently sort and limit the result set to the ten most recent sales.
* B (Incorrect): Replacing FETCH FIRST with ROWNUM would not necessarily improve the performance unless there is an appropriate index that the optimizer can use to avoid sorting the entire result set.
* C (Incorrect): There is no indication that the current statistics are inaccurate; hence, collecting new statistics may not lead to performance improvement.
* D (Incorrect): While adaptive plans can provide performance benefits by allowing the optimizer to adapt the execution strategy, the main issue here is the lack of an index on the ORDER BY column.
* E (Incorrect): Creating an index on SALES.CUST_ID could improve join performance but would not address the performance issue caused by the lack of an index on the ORDER BY column.
References:
* Oracle Database SQL Tuning Guide: Managing Indexes
* Oracle Database SQL Tuning Guide: Using Indexes and Clusters

NEW QUESTION # 36
You must write a statement that returns the ten most recent sales. Examine this statement:

```
SELECT s.prod_id, p.prod_name, s.amount_sold, c.cust_name, c.cust_city
FROM   sales s, customers c, products p
WHERE  s.cust_id = c.cust_id AND
       s.prod_id = p.prod_id
ORDER BY s.time_id
FETCH FIRST 10 ROWS ONLY;
```

Users complain that the query executes too slowly. Examine the statement's current execution plan:

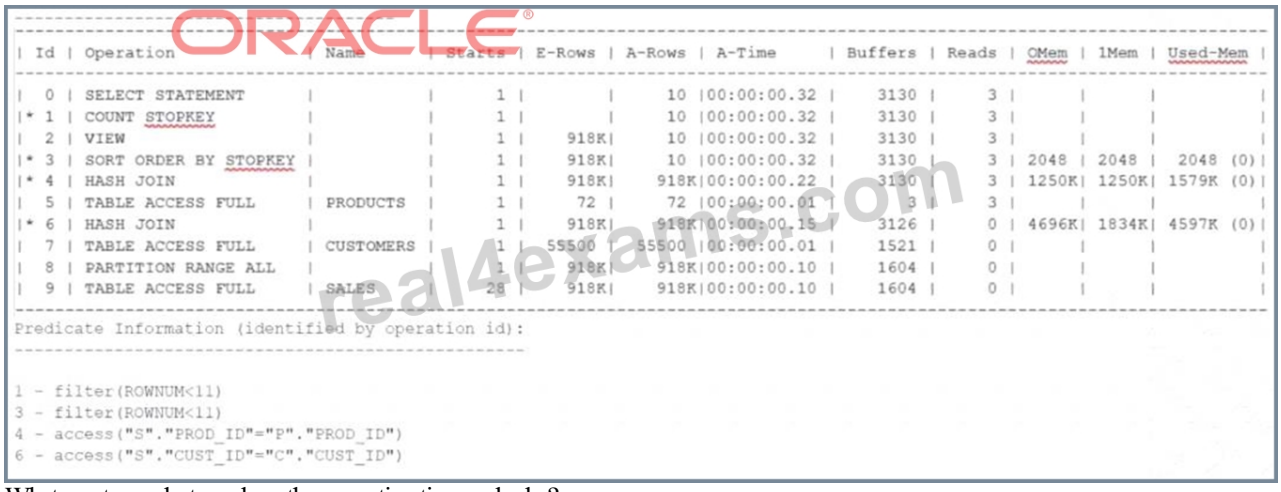

```
| Id | Operation                | Name      | Starts | E-Rows | A-Rows | A-Time      | Buffers | Reads | OMem | 1Mem | Used-Mem |
|  0 | SELECT STATEMENT         |           |      1 |        |     10 |00:00:00.32 |    3130 |     3 |      |      |          |
|* 1 | COUNT STOPKEY            |           |      1 |        |     10 |00:00:00.32 |    3130 |     3 |      |      |          |
|  2 | VIEW                     |           |      1 |   918K|     10 |00:00:00.32 |    3130 |     3 |      |      |          |
|* 3 | SORT ORDER BY STOPKEY    |           |      1 |   918K|     10 |00:00:00.32 |    3130 |     3 | 2048 | 2048 |  2048 (0)|
|* 4 | HASH JOIN                |           |      1 |   918K|   918K|00:00:00.22 |    3130 |     3 | 1250K| 1250K| 1579K (0)|
|  5 | TABLE ACCESS FULL        | PRODUCTS  |      1 |    72 |     72 |00:00:00.01 |       3 |     3 |      |      |          |
|* 6 | HASH JOIN                |           |      1 |   918K|   918K|00:00:00.15 |    3126 |     0 | 4696K| 1834K| 4597K (0)|
|  7 | TABLE ACCESS FULL        | CUSTOMERS |      1 |  55500 |  55500 |00:00:00.01 |    1521 |     0 |      |      |          |
|  8 | PARTITION RANGE ALL      |           |      1 |   918K|   918K|00:00:00.10 |    1604 |     0 |      |      |          |
|  9 | TABLE ACCESS FULL        | SALES     |     28 |   918K|   918K|00:00:00.10 |    1604 |     0 |      |      |          |

Predicate Information (identified by operation id):

1 - filter(ROWNUM<11)
3 - filter(ROWNUM<11)
4 - access("S"."PROD_ID"="P"."PROD_ID")
6 - access("S"."CUST_ID"="C"."CUST_ID")
```

What must you do to reduce the execution time and why?

- A. Create an index on SALES.CUST_ID to force an INDEX RANGE SCAN on this index followed by a NESTED LOOP join between CUSTOMERS and SALES.
- B. Replace the FETCH FIRST clause with ROWNUM to enable the use of an index on SALES.
- C. Create an index on SALES.TIME_ID to force the return of rows in the order specified by the ORDER BY clause.
- D. Collect a new set of statistics on PRODUCT, CUSTOMERS, and SALES because the current stats are inaccurate.
- E. Enable Adaptive Plans so that Oracle can change the Join method as well as the Join order for this query.

Answer: C

Explanation:
The execution plan shows a full table access for theSALEStable. To reduce the execution time, creating an index onSALES.TIME_IDwould be beneficial as it would allow the database to quickly sort and retrieve the most recent sales without the need to perform a full table scan, which is I/O intensive and slower. By indexing TIME_ID, which is used in theORDER BYclause, the optimizer can take advantage of the index to efficiently sort and limit the result set to the ten most recent sales.
* B (Incorrect):ReplacingFETCH FIRSTwithROWNUMwould not necessarily improve the performance unless there is an appropriate index that the optimizer can use to avoid sorting the entire result set.
* C (Incorrect):There is no indication that the current statistics are inaccurate; hence, collecting new statistics may not lead to performance improvement.
* D (Incorrect):While adaptive plans can provide performance benefits by allowing the optimizer to adapt the execution strategy, the main issue here is the lack of an index on theORDER BYcolumn.
* E (Incorrect):Creating an index onSALES.CUST_IDcould improve join performance but would not address the performance issue caused by the lack of an index on theORDER BYcolumn.
References:
* Oracle Database SQL Tuning Guide:Managing Indexes
* Oracle Database SQL Tuning Guide:Using Indexes and Clusters

NEW QUESTION # 37
Which three statements are true about using the in Memory (IM) column store?

- A. It does not require all database data to fit in memory to improve query performance.
- B. It can improve OLTP workload performance by avoiding the use of indexes.
- C. It does not improve performance for queries using cached results of function evaluations on columns from the same table.
- D. It does not improve performance for queries using user-defined virtual column results.
- E. It improves performance for queries joining several tables using bloom filter joins.
- F. It does not improve performance for queries that use join groups on columns from different tables.

Answer: A,B,E

Explanation:
The Oracle In-Memory (IM) column store feature enhances the performance of databases by providing a fast columnar storage format for analytical workloads while also potentially benefiting OLTP workloads.
* C (True):It can improve OLTP workload performance by providing a faster access path for full table scans and reducing the need for indexes in certain scenarios, as the In-Memory store allows for efficient in-memory scans.

* E (True):The In-Memory column store does not require all database data to fit in memory. It can be used selectively for performance-critical tables or partitions, and Oracle Database will manage the population and eviction of data as needed.
* F (True):In-Memory column store can significantly improve performance for queries joining several tables, especially when bloom filters are used, as they are highly efficient with the columnar format for large scans and join processing.
The other options provided are not correct in the context of the In-Memory column store:
* A (False):While In-Memory column store is designed for analytical queries rather than caching results of function evaluations, it does not specifically avoid improving performance for queries using cached results of function evaluations.
* B (False):In-Memory column store can improve the performance of queries that use join groups, which can be used to optimize joins on columns from different tables.
* D (False):In-Memory column store can improve the performance of queries using expressions, including user-defined virtual columns, because it supports expression statistics which help in
* optimizing such queries.
References:
* Oracle Database In-Memory Guide:In-Memory Column Store in Oracle Database
* Oracle Database In-Memory Guide:In-Memory Joins
* Oracle Database In-Memory Guide:In-Memory Aggregation

NEW QUESTION # 38
Which two statements are true about Data Pump import for objects that used the in Memory (IM) column store in their source database?

- A. It must always transports existing INMEMORY attributes.
- B. Its TRANSFORM clause can be used to add the INMEMORV clause to exported tables that lack them.
- C. It always gives preference to the IM column store clause defined at the tablespace level over table-level definitions.
- D. Its INMEM0RY_CLAUSE of the Data Pump Export allows modifications to IM column store clause of a table with existing INMEMORY setting.
- E. It can generates the INMEMORY clause that matches the table settings at export time.
- F. It ignores the IM column store clause of the exporting objects.

Answer: B,E

Explanation:
When importing objects that used the In-Memory (IM) column store in their source database using Oracle Data Pump, the following statements are true:
* D (Correct):TheTRANSFORMclause can be used to alter object creation DDL during import operations. This can include adding theINMEMORYclause to tables that were not originally using the IM column store.
* F (Correct):The import operation can preserve theINMEMORYattributes of tables as they were at the time of export, effectively replicating the IM column store settings from the source database.
The other statements are not accurate in the context of Data Pump import:
* A (Incorrect):Data Pump does not give preference to the IM column store clauses at the tablespace level over table-level definitions unless explicitly specified by theTRANSFORMclause.
* B (Incorrect):While Data Pump can transport existingINMEMORYattributes, it is not mandatory. It is controlled by theINCLUDEorEXCLUDEData Pump parameters or theTRANSFORMclause.
* C (Incorrect):TheINMEMORY_CLAUSEparameter is not part of the Data Pump Export utility. To modify the IM column store clauses, you would use theTRANSFORMparameter during import, not export.
* E (Incorrect):Data Pump does not ignore the IM column store clause unless specifically instructed to do so via theEXCLUDEparameter.
References:
* Oracle Database Utilities:Data Pump Export
* Oracle Database Utilities:Data Pump Import

NEW QUESTION # 39
.....

Knowledge about a person and is indispensable in recruitment. That is to say, for those who are without good educational background, only by paying efforts to get an acknowledged 1z1-084 certification, can they become popular employees. So for you, the 1z1-084 latest braindumps complied by our company can offer you the best help. With our test-oriented 1z1-084 Test Prep in hand, we guarantee that you can pass the 1z1-084 exam as easy as blowing away the dust, as long as you guarantee 20 to 30 hours

practice with our 1z1-084 study materials. The reason why we are so confident lies in the sophisticated expert group and technical team we have, which do duty for our solid support.

Latest 1z1-084 Exam Price: https://www.real4exams.com/1z1-084_braindumps.html

- Reliable 1z1-084 Test Forum □ Valid Braindumps 1z1-084 Ebook □ 1z1-084 Reliable Exam Answers □ Go to website { www.testsdumps.com } open and search for ➠ 1z1-084 □ to download for free □1z1-084 Latest Exam Tips
- Professional 1z1-084 Exam Outline - Trusted Latest 1z1-084 Exam Price - New 1z1-084 Official Practice Test □ Go to website □ www.pdfvce.com □ open and search for [1z1-084] to download for free □Printable 1z1-084 PDF
- 100% Pass Quiz 2025 Oracle 1z1-084: Reliable Oracle Database 19c Performance and Tuning Management Exam Outline □ Search for ▹ 1z1-084 ◃ and download it for free on ☀ www.torrentvce.com □☀□ website □1z1-084 Interactive EBook
- Free PDF 1z1-084 - Oracle Database 19c Performance and Tuning Management High Hit-Rate Exam Outline □ Search for ▹ 1z1-084 ◃ and download it for free immediately on 【 www.pdfvce.com 】 □1z1-084 Testking Learning Materials
- Obtained The 1z1-084 Exam Outline Means Passing The Oracle Database 19c Performance and Tuning Management □ Search for ⇛ 1z1-084 ⇚ and easily obtain a free download on ➡ www.examdiscuss.com □□□ □Complete 1z1-084 Exam Dumps
- Exam 1z1-084 Fees □ Reliable 1z1-084 Test Forum □ Reliable 1z1-084 Dumps Sheet □ Open website 「 www.pdfvce.com 」 and search for ⇛ 1z1-084 ⇚ for free download □Authorized 1z1-084 Test Dumps
- Reliable 1z1-084 Dumps Sheet □ 1z1-084 Latest Exam Tips □ Exam 1z1-084 Fees □ Simply search for □ 1z1-084 □ for free download on （ www.real4dumps.com ） □Exam 1z1-084 Fees
- Authorized 1z1-084 Test Dumps □ Authorized 1z1-084 Test Dumps □ 1z1-084 Exam Dumps.zip □ Enter ➡ www.pdfvce.com □□□ and search for ➡ 1z1-084 □□□ to download for free □Real 1z1-084 Exam Answers
- Real 1z1-084 Exam Answers □ 1z1-084 Exam Dumps.zip ☃ Valid Braindumps 1z1-084 Ebook □ The page for free download of ☀ 1z1-084 □☀□ on ➤ www.examcollectionpass.com □ will open immediately □Reliable Study 1z1-084 Questions
- Professional 1z1-084 Exam Outline - Trusted Latest 1z1-084 Exam Price - New 1z1-084 Official Practice Test □ Open ⇛ www.pdfvce.com ⇚ and search for ▹ 1z1-084 ◃ to download exam materials for free □Valid Braindumps 1z1-084 Ebook
- Pass Guaranteed Quiz Oracle - 1z1-084 –Valid Exam Outline □ 「 www.examdiscuss.com 」 is best website to obtain ▸ 1z1-084 ◂ for free download □1z1-084 Exam Dumps.zip
- www.myaniway.com, codifyedu.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, motionentrance.edu.np, study.stcs.edu.np, www.wcs.edu.eu, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, pct.edu.pk, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

DOWNLOAD the newest Real4exams 1z1-084 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=12zmuODnh_5V_atFMNTiAKgC5-rqy-Zr2