

# Can we have more hot, please?

How we achieved sustainable growth for powerful hot data, using SAP HANA extension nodes and the identification of tiering potential by Oxygen for HANA. A compelling customer case by ONE-Consultants.

## Robust platform

For the past 5 years, the HANA platform in question is used to report on a retail dataset with billions of records. Dashboards as well as detailed reports are being powered by views. Over time, the dataset has been growing and has been extended to allow for more insights.

The HANA database and modeling environment allows for very fast querying and data processing on large volumes, agile delivery of functionality through views, virtual data consumption and the integration of various datasets. As the dataset is growing and other areas and business units are onboarding to the platform, ONE needed to find a solution to handle this projected growth in a sustainable way. What data could be off-loaded to warm or cold storage solutions without compromising business value and user acceptance?

## Robust Extension

As a data tiering solution, SAP HANA Extension Node offered the perfect fit:

1. Allows similar SAP HANA performance
2. Has no adverse effect on the querying scenarios
3. Needs no modeling changes
4. As a native HANA solution, it allows quick implementation

A short note on the extension node concept: a normal HANA node's memory is occupied with 50% data only, reserving 50% of memory space for the calculation of (sub)results for reporting and loading. An extension node doesn't have this 50% reservation constraint and can even have a 200% data allocation. This means that the system will swap data in and out the extension memory node depending on the need. The traffic on this node should be lower as it is intended to use for less frequent reported or processed datasets.

With regards to point 1 above, our tests concluded no significant drop in reporting performance<sup>1</sup> even with -worst case- a fully loaded extension node where the requested dataset had to be retrieved into memory first. With regards to point 2, alternative warm solution attempts were found to have adverse effect on exception calculation scenarios amongst other, rendering the solution not fit. An extension node being a natural part of the HANA platform has no such limitation. As stated in the third point, no modeling changes are needed. This helps to keep our models lean and supportable, not compromising our agile project methodology.

## Robust Insights

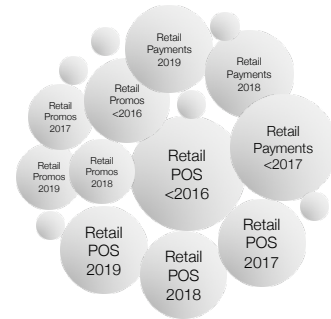
An insight into the tables' content by year for this Retail area was necessary. Oxygen<sup>2</sup> provided us that insight and was also able to uncover less-known tables - a typical consequence of organic development. As the business confirmed most of the reporting is done for 2 rolling years, we decided to keep 2019 till 2017 hot and moved the older years to the warm extension node and even older data to cold store, after re-partitioning. Similar for complete tables less frequently reported on. Tables that were identified to be staging-use-only were kept in hot or moved to warm. Even the deletion of data was considered and executed.

---

<sup>1</sup> In our case, the table(s) were already allocated across nodes, so inter-node communication was already occurring.

<sup>2</sup> Oxygen is a SAP HANA application providing insight in the HANA dataset by evaluating the in-memory data and interpreting their purpose.

**“The insights were a catalyst to think and decide upon the proper use, purpose and upkeep of the whole dataset.”**



### Robust Numbers

The numbers in graph 1 show the results of the data size in Gb per data tiering execution stage in each tier. (November and December activities are planned for at the time of this article).

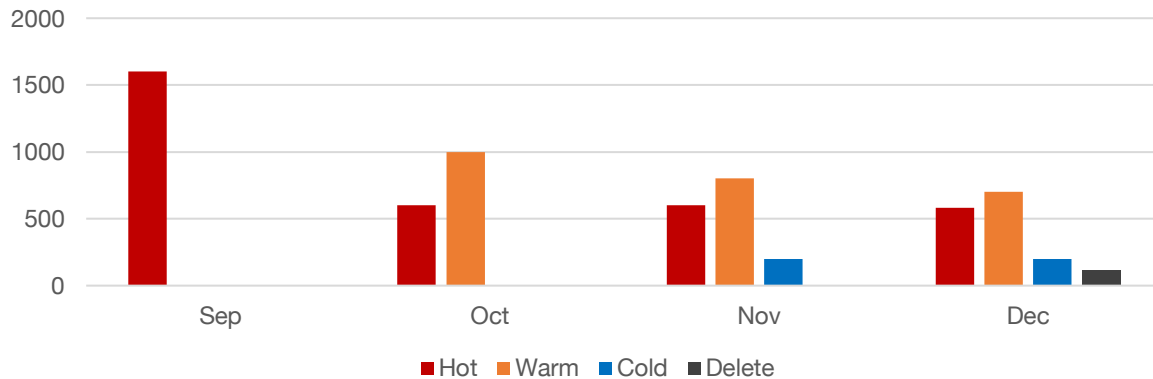


Figure 1 Data Tiering Status per Month in Gb.

As a result, we can now continue to expand and at the same time access the older data as and when needed. Room – Oxygen – for sustainable growth is created.

**“This exercise created room for additional datasets in hot memory for fast processing and advanced analytic purposes.”**

### Impact on TCO

Let’s assume that we only have this one dataset on the HANA platform. Using 512Gb memory nodes, we needed  $1600\text{Gb} \times 2 = 3200/512\text{Gb} = 7$  nodes in the all-hot situation. After the first and second data tiering execution inspired by Oxygen insights, we now need:  $600\text{Gb} \times 2 = 1200/512\text{Gb} = 3$  nodes plus 1 extension node (800Gb of warm data fits on a 512Gb extension node @ 200% allocation, with warm room left). The cost saving opportunity consists of going from 7 to 4 nodes, i.e. 43%. Or the opportunity to repurpose the resulting 3 empty hot nodes for new business value.

### About Oxygen

Oxygen is a HANA application that evaluates and interprets the in-memory data. It is available at ONE-Labs<sup>3</sup>. The case above describes the use and results for 1 business area. Oxygen reports on all business areas on a SAP HANA system. Hence, the above exercise can be done for the whole platform, unlocking value, by making room on a budgetary or functional level.

More info on Oxygen and ONE-Labs can be found at: [www.one-labs.com/oxygen](http://www.one-labs.com/oxygen)

<sup>3</sup> ONE-Labs produces solutions for the SAP HANA platform using the ideas, expertise and experience of partner consulting company ONE-Consultants. ONE-Consultants was founded in 2008 and provides SAP HANA related services (strategy, modeling, data sourcing, performance advisory, security, system administration...) at major global customers.