

THE TRUE STORY OF A VERY BIG CMDB



A large customer approached BMC recently with the requirement to store BILLIONS of items in a CMDB-like construct. This customer had already acquired the BMC TrueSight solution and wanted an integrated approach.

When we considered the customer use case, one thing became clear: This request would result in what we believe is the biggest CMDB currently deployed in any organization, anywhere in the world—road-mapped to a billion “things” by 2020 and continuing to scale beyond that massive number.

This was not your average CMDB, and it was not an average CMDB implementation, either. The key here was scale and the diversity of data the customer needed to store, including significant non-IT and IoT-related data.

The key to this implementation lay in the flexibility of BMC Helix CMDB itself. Although we supply a predetermined default data model (the CDM), this doesn't mean a customer has to use the default if it doesn't fit their business needs. BMC Helix CMDB is specifically designed to allow customers to tailor the data being held either within the CDM or to create their own hierarchy to store their data the way they choose.

Our customer designed a classification and hierarchy to meet their business needs, based on their specific requirements.

Then they tackled the task of identifying from where the data needed to come. This step is a critical

part of any CMDB implementation, whether it's 100,000 or one billion "things." You must understand what data you *need* and then *where* it will come from; we don't just populate a CMDB with data without understanding *why* we are doing it.

Due to the complex nature of the customer's source data, we used the power of Atrium Integrator and its underlying technology to assemble, format, and shape the data using transformations ready for import to the CMDB dataset structure.

Once the data population task was complete, we turned our focus to the aims of the users. How do we ensure responsive searches? How do we ensure users, both human and integration, can easily identify the data they need? The customer's partnership with BMC Platinum Engineering resulted in product optimizations being identified and implemented to drive the consumption performance needs.

Additionally, this engagement resulted in performance tests with the large sample data, and replication in a customer-equivalent test environment to prove out that the underlying application could meet the goals and requirements laid down by the customer's architecture group.

Make room for the Enterprise of Things

This particular customer's needs might be extraordinary, but other organizations will eventually need to cope with greater demands placed on their CMDB. This will be especially true as use cases expand for the Internet of Things (IoT)—or Enterprise of Things, a phrase we like for describing how businesses embrace IoT. CMDBs need to be prepared for handling more "things" in the future, and BMC Helix CMDB has already proven capable of handling the *most* things.

As noted with this particularly large, complex project, you do need to optimize CMDB implementations for their intended use cases, and there can be challenges at scale. That's true even if your organization has far fewer than a billion "things" to handle. There are key issues to keep in mind during planning as you shift from a straightforward, traditional infrastructure to a more complex Enterprise of Things environment.

Many of these issues stem from usability and performance—issues that anyone who has dealt with large databases can understand.

Massive CMDB? No sweat.

BMC is ready for the challenge, having already helped customers like the one above with the most complex implementation and optimization challenges. We've dealt with them first-hand and can help customers plan for and solve problems as needed to meet their particular needs.

We've developed best practices and recommendations for ensuring that performance and usability don't suffer in a large-scale CMDB. CI search responsiveness is a key area to measure and optimize, for instance, as is speed to populate. Database tuning is critical for consistent performance over time, and application tuning can lead to further performance gains. Search performance varies between the operators, and we're able to change the CMDB UI to ensure the optimal operator for search performance is set as the default. And these are just the highlights.

BMC doesn't see large CMDBs as problems. We've already got the solutions.

Oh, and did we mention that this customer's very large CMDB runs entirely in the cloud on AWS, and

not in their on-premises infrastructure? Some would say that can't be done. We went ahead and did it anyway.

Our customer came to us with, literally, the biggest CMDB needs, and we worked together to meet those needs. As the Enterprise of Things expands, so too will the demands placed on an organization's CMDB. As we've learned firsthand, large CMDBs aren't really problems—they just require some extra planning and configuration.

[BMC Helix CMDB](#) is up for the challenge.

BMC Helix CMDB can handle the biggest of use cases, no problem.

We can even help run it entirely in the cloud.

Learn more about [BMC Helix CMDB](#) today.