

Interactive | Analytics

An Origin Story

From SCUBA to Inter|ana

The core of what we have built at Interana was born while our founders were at Facebook. At Facebook, they realized that they needed to look at how users were behaving and interacting with the platform in a much more granular way than was possible with any commercially available products. Like many web-scale companies then, they chose to invest heavily in Hadoop. Like many that have since found out, while Hadoop is great for low cost storage, it could not deliver on the need to be able to interact with and explore that data.

When we developed what is now SCUBA, we went beyond Hadoop, and the past 20 years of DW/BI solutions and were able to make the analytics fast and visual. What is different about fast and visual is that it gives you the ability to do real-time discovery and interactivity. This has been the long-time holy grail. After they perfected this for Facebook, our founders chose to leave and build their own company around the know-how they then had and built a 2.0 version for the Enterprise. **Interana is that platform.**

To deliver that real-time discovery and interactivity, you need a platform that is simultaneously extremely fast across the full dataset, and an interface that is intuitive and visual. This requires a dramatically different architecture. Interana's architecture is different across the whole stack, from the UI controls, to what primitives the engine computes, to how the data is stored, to how it is pipelined from disks through CPU caches. We reimagined each of these in a way that puts the time dimension at the forefront, and values user interactivity. This gives an elegance in the design that makes each layer simple and fast. We make it easy for our users to understand what's going on, and fast for the platform to execute queries. And you must have both things to realize value.

We know that changing one layer of the stack at a time just doesn't work. You could build something that looks like what Interana does, but it's at least two orders of magnitude slower or smaller. It will also require mental gymnastics for users to understand a model that isn't consistent all the way through the stack. 100x times slower doesn't just mean it feels sluggish, it means you use it in a fundamentally different way that is much less valuable. Because we have the real magic in building our interface and engine to work together, with a model that is simple to understand and fast to execute.

With Interana, something that was previously impossible, is now fast and easy to deploy and use.

To deliver that real-time discovery and interactivity, you need a platform that is simultaneously extremely fast across the full dataset, and an interface that is intuitive and visual.

Bobby Johnson
Founder and CTO