In this talk I present our recent result, the first scalable distributed storage system that is provably robust against massive storage failures, even when they are caused by an insider, i.e., someone who has complete knowledge of the storage system.

More specifically, we developed a method to distribute and encode data in a storage system consisting of n servers with just constant redundancy so that for any set of $\gamma n^{1/\log\log n}$ servers with corrupted storage information, where $\gamma>0$ is constant, the system can serve any set of lookup and write requests, one per server, in polylogarithmic time and with polylogarithmic work per server.

Previously, scalable solutions for attacks of that scale were only known for crash failures, i.e., servers simply become unavailable.

This is a much simpler case since it is a priori clear which servers are problematic and which are not, and problematic servers stay silent.

The key idea behind our new construction is to combine techniques from authenticated data structures, distributed coding and shared memory emulations.

Since our authentication mechanism is based on Merkle hash trees, which require the existence of one-way hash functions to work, the only limitation that we have on the storage corruption is that it is due to a polynomially bounded insider.