



To our shareowners:

Random forests, naïve Bayesian estimators, RESTful services, gossip protocols, eventual consistency, data sharding, anti-entropy, Byzantine quorum, erasure coding, vector clocks ... walk into certain Amazon meetings, and you may momentarily think you've stumbled into a computer science lecture.

Look inside a current textbook on software architecture, and you'll find few patterns that we don't apply at Amazon. We use high-performance transactions systems, complex rendering and object caching, workflow and queuing systems, business intelligence and data analytics, machine learning and pattern recognition, neural networks and probabilistic decision making, and a wide variety of other techniques. And while many of our systems are based on the latest in computer science research, this often hasn't been sufficient: our architects and engineers have had to advance research in directions that no academic had yet taken. Many of the problems we face have no textbook solutions, and so we -- happily -- invent new approaches.

Our technologies are almost exclusively implemented as *services*: bits of logic that encapsulate the data they operate on and provide hardened interfaces as the only way to access their functionality. This approach reduces side effects and allows services to evolve at their own pace without impacting the other components of the overall system. Service-oriented architecture -- or SOA -- is the fundamental building abstraction for Amazon technologies. Thanks to a thoughtful and far-sighted team of engineers and architects, this approach was applied at Amazon long before SOA became a buzzword in the industry. Our e-commerce platform is composed of a federation of hundreds of software services that work in concert to deliver functionality ranging from recommendations to order fulfillment to inventory tracking. For example, to construct a product detail page for a customer visiting Amazon.com, our software calls on between 200 and 300 services to present a highly personalized experience for that customer.

State management is the heart of any system that needs to grow to very large size. Many years ago, Amazon's requirements reached a point where many of our systems could no longer be served by any commercial solution: our key data services store many petabytes of data and handle millions of requests per second. To meet these demanding and unusual requirements, we've developed several alternative, purpose-built persistence solutions, including our own key-value store and single table store. To do so, we've leaned heavily on the core principles from the distributed systems and database research communities and invented from there. The storage systems we've pioneered demonstrate extreme scalability while maintaining tight control over performance, availability, and cost. To achieve their ultra-scale properties these systems take a novel approach to data update management: by relaxing the synchronization requirements of updates that need to be disseminated to large numbers of replicas, these systems are able to survive under the harshest performance and availability conditions. These implementations are based on the concept of eventual consistency. The advances in data management developed by Amazon engineers have been the starting point for the architectures underneath the cloud storage and data management services offered by Amazon Web Services (AWS). For example, our Simple Storage Service, Elastic Block Store, and SimpleDB all derive their basic architecture from unique Amazon technologies.

Other areas of Amazon's business face similarly complex data processing and decision problems, such as product data ingestion and categorization, demand forecasting, inventory allocation, and fraud detection. Rule-based systems can be used successfully, but they can be hard to maintain and can become brittle over time. In many cases, advanced machine learning techniques provide more accurate classification and can self-heal to adapt to changing conditions. For example, our search engine employs data mining and machine learning algorithms that run in the background to build topic models, and we apply information extraction algorithms to identify attributes and extract entities from unstructured descriptions, allowing customers to narrow their searches and quickly find the desired product. We consider a large number of factors in search relevance to

predict the probability of a customer's interest and optimize the ranking of results. The diversity of products demands that we employ modern regression techniques like trained random forests of decision trees to flexibly incorporate thousands of product attributes at rank time. The end result of all this behind-the-scenes software? Fast, accurate search results that help you find what you want.

All the effort we put into technology might not matter that much if we kept technology off to the side in some sort of R&D department, but we don't take that approach. Technology infuses all of our teams, all of our processes, our decision-making, and our approach to innovation in each of our businesses. It is deeply integrated into everything we do.

One example is Whispersync, our Kindle service designed to ensure that everywhere you go, no matter what devices you have with you, you can access your reading library and all of your highlights, notes, and bookmarks, all in sync across your Kindle devices and mobile apps. The technical challenge is making this a reality for millions of Kindle owners, with hundreds of millions of books, and hundreds of device types, living in over 100 countries around the world—at 24x7 reliability. At the heart of Whispersync is an eventually consistent replicated data store, with application defined conflict resolution that must and can deal with device isolation lasting weeks or longer. As a Kindle customer, of course, we hide all this technology from you. So when you open your Kindle, it's in sync and on the right page. To paraphrase Arthur C. Clarke, like any sufficiently advanced technology, it's indistinguishable from magic.

Now, if the eyes of some shareowners dutifully reading this letter are by this point glazing over, I will awaken you by pointing out that, in my opinion, these techniques are not idly pursued – they lead directly to free cash flow.

We live in an era of extraordinary increases in available bandwidth, disk space, and processing power, all of which continue to get cheap fast. We have on our team some of the most sophisticated technologists in the world – helping to solve challenges that are right on the edge of what's possible today. As I've discussed many times before, we have unshakeable conviction that the long-term interests of shareowners are perfectly aligned with the interests of customers.

And we like it that way. Invention is in our DNA and technology is the fundamental tool we wield to evolve and improve every aspect of the experience we provide our customers. We still have a lot to learn, and I expect and hope we'll continue to have so much fun learning it. I take great pride in being part of this team.

As always, I attach a copy of our original 1997 letter. Our approach remains the same, and it's still Day 1.

A handwritten signature in black ink, reading "Jeff P. Bezos". The signature is fluid and cursive, with a long horizontal stroke at the end.

Jeffrey P. Bezos  
Founder and Chief Executive Officer  
Amazon.com, Inc.