

# Ad Tech Supply Chain Carbon Footprint

*A research project determining the true carbon footprint of the digital ad industry and proposing a cleaner alternative.*

The digital ad supply chain is among the most wasteful ones. The reason is the compounding effect of multiple inefficiencies routed deeply in the underlying infrastructure: from inefficient bidding, to the necessity to store the same databases on multiple instances for the different stakeholders, to poor code and file-size optimisations. Given that the supply chain stores tens of billions of events per day, those quickly add up.

Research is already being done in this space. A [recent study by Fifty-Five](#), a single digital campaign can generate more than 70 tons of CO<sub>2</sub>e, the equivalent of the carbon footprint of about seven people for a year.

[Scope3](#), a company founded this year by the inventor of Real-Time Bidding (RTB) advertising, Brian O'Kelley, is focusing specifically on this problem. We are in touch with mister O'Kelley and his company is interested in collaborating on this research project. According to Scope3 in the U.K., France, and Australia, it's the number one topic and there is immense interest from agencies and brands. In the U.S., it's number two, maybe number three in priority. DEI is probably the hottest topic in the U.S. Plus cost-control and some of the Big Tech issues are concerns, but sustainability issues are definitely rising.

AdHash is interested in furthering the research on this topic because for the past four years, we have been busy reengineering the ad supply chain from the ground up. We did so by developing a much more transparent protocol for bidding. We then created the platforms for advertisers (demand-side) and publishers (supply-side), as well as the bidder that does the market making, and the interface for end users to express their preferences.

What makes our approach fundamentally different is that we eliminate the need for intermediaries on the supply chain. By creating a distributed architecture for ad serving, we give advertisers complete control over their targeting and reporting, thus eliminating the need to rely and up to 25 different categories of middlemen just to serve a single ad. This should in theory reduce the carbon footprint of the ad supply chain by orders of magnitude. Here is how:

Let's take Google for an example. When they manage the data for 1,000 clients, they must operate one database that is 1,000 times the size of each individual client's. AdHash, on the other hand, allowing advertisers and publishers to self-host, results in 1,000 databases that are each 1/1,000-th of the size of Google's. It becomes exponentially more expensive to manage a database as it grows. So a database that is 1,000 times bigger might cost 5,000 or 10,000 more in terms of resources. This is a fairly straightforward part of the project and could be emulated on AWS or another cloud providers with the help of our development team.

Another interesting aspect is the elimination of duplicated requests. Having unique ad IDs (we call them ad hashes) means that the same advertiser, even if coming from multiple exchanges, only has to bid once. So that reduces further the number of server requests.

The elimination of [ad fraud](#) before it even happens also has a significant effect. It reduces the carbon footprint of data associated with it that have to be stored. Many brands and agencies use multiple verification companies (MOAT, Integral Ad Science, DoubleVerify) to tell them that 20-30% of their traffic is invalid at the end of the month. That means their DSP, exchange, ad server, etc. all have to store those additional 20-30%. And the verification company has to store the data about the fraud as well. AdHash doesn't need to do that. We just store the number of impressions or clicks associated with fraud or invalid activity and the reason why we discounted them.

Eliminating ad fraud also means that all the server farms engaged in perpetrating it would be stopped.

Probably the lowest-hanging fruit in this research is the fact that the current ad supply chain requires 25-30 different middlemen to store the same data; it's 25-30 companies storing their version of the same data. When using AdHash, only the publisher and advertiser store the data and the AdHash Bidder stores a much lighter version of it to serve as a reference point. In addition to this, advertisers and publishers can actually opt-in to not store any impression data, just click data, reducing the required storage by two additional orders of magnitude.

Another easy target to optimise is the ad file size. Our company, for example, only allows static and very light ads. A typical AdHash ad is 100-150 kB. A typical Google ad is probably 1-3 MB on average. Some rich-media formats can even exceed 10 MB.

We are interested in working with ambitious students who would spend a few months investigating these and other aspects of the carbon footprint of the digital ad supply chain. The goal is not to create a PR publication for AdHash, but rather to truly estimate the cost of the existing infrastructure. We are providing AdHash as a point of reference that shows how things could be done better today. There surely exist even better ideas how to do it and we would be thrilled to learn about them and include them in the study. The students can count on our team's full support in both research and development. We would also be happy to provide support in future career development if this topic becomes something the students want to be involved with long-term.