

Blank, on purpose.

Score each vendor 0, 1, or 2. Zero if they cannot answer, one if they answer but cannot show you, two if they answer and can demonstrate it on your data. Add the columns. The total matters less than which questions produced the zeros, because the zeros are the ones you will be living with.

VENDOR A NAME		VENDOR B NAME		VENDOR C NAME	
#	QUESTION	VENDOR A	VENDOR B	VENDOR C	
1	Is This Number Causal or Correlational? THE METHOD				
2	Does It Run Experiments, or Only Observe? THE METHOD				
3	Does the Specification Exclude Mediators? THE METHOD				
4	Can It Connect Spend to the Transaction That Actually Matters? THE DATA				
5	Can It Operate Without Exposing PHI, and Will It Sign a BAA? THE DATA				
6	Who Owns the Tracking Layer, and Who Is Watching It? THE DATA				
7	Does It Publish Uncertainty Intervals, and Are They Used? THE HONESTY				
8	When Two Methods Disagree, Does It Reconcile Them or Pick One? THE HONESTY				
9	Will It Tell You When It Cannot Answer? THE HONESTY				
10	Will It Validate Against a Period Where You Already Know the Answer? THE PROOF				
11	Does It Model Your Locations Jointly or One at a Time? THE PROOF				
12	Is Measurement Independent of Media Execution? THE PROOF				
Total (out of 24)					

What a strong and a weak answer sound like, for each question, is at mediaura.io/evaluation-framework. Reuse this sheet in an RFP, a vendor scorecard, or a board memo without asking. If it is useful, a link back is appreciated but never required.