



SEPTEMBER 2026 EDITION

STANDARD EDITION PACK
REVENUE OPERATIONS

CRM

What twelve AI models recommend when a buyer asks for CRM.

12 models at every segment

The standard tier: each lab's low-cost model, asked on behalf of each of the three buyers.

140 first choices

Across the direct, paraphrase, budget and scale prompts, all segments.

1,202 judge labels

One per product named, each with its verbatim evidence.

Six framings per segment

Direct, paraphrase, comparative, budget, scale and negative.

WHICH MODELS, FOR WHICH BUYER

Twelve models answered for every buyer segment, the standard tier of the public index. Its standing here is the standing the site publishes.

HOW IT WAS ASKED

Six framings of the same buying question, on behalf of a small B2B company, a mid-market B2B company and an enterprise B2B company. Every figure is derived from the raw judge labels under vendor table v2026-09-15.2. Nothing here is a recommendation by the index.


Contents

Each buyer segment carries the same four sections in the same order. The standing is the answer; what they warned about and what they cite are the evidence; how they answered is the full record for the segment.

Front matter	Cover, contents, executive summary	1
—	Executive summary	3
Small business	Asked of the standard twelve models	4
—	The standing	4
—	What they warned about	5
—	What they cite	6
—	How they answered	7
Mid-market	Asked of the standard twelve models	10
—	The standing	10
—	What they warned about	11
—	What they cite	12
—	How they answered	13
Enterprise	Asked of the standard twelve models	16
—	The standing	16
—	What they warned about	17
—	What they cite	18
—	How they answered	19
Method, in brief	The run, the judge, the noise floor and the vendor table	22

HubSpot CRM leads at small business and mid-market; Salesforce Sales Cloud leads at enterprise. The leader changes once across the segments. This is the first edition of the standard tier, so no movement can be computed; from the next edition each standing carries its move against the noise floor of 10 points. Every figure in this report is a row in standing.csv, labels.csv, responses.csv, prompts.csv, models.csv and index.sqlite, with the data dictionary in README.md.

Executive summary



The leader changes once across the segments: HubSpot CRM holds small business and mid-market, and Salesforce Sales Cloud leads at enterprise, where HubSpot CRM is itself argued against.

	SMB	MID-MARKET	ENTERPRISE
WARNED AGAINST	Salesforce Sales Cloud, 76% of 34 labels. Microsoft Dynamics 365 Sales, 82% of 11 labels.	Salesforce Sales Cloud, 35% of 49 labels. HubSpot CRM, 19% of 54 labels.	Pipedrive, 75% of 16 labels. Salesforce Sales Cloud, 19% of 62 labels.
HELD UNDER REWORDING	Four of twelve held: Grok 4.1 Fast, Llama 4 Maverick, GLM 4.7 FlashX, MiniMax M2.5.	Five of twelve held: Claude Haiku 4.5, DeepSeek V4 Flash, Llama 4 Maverick, Kimi K2 and one more.	Eight of twelve held: Claude Haiku 4.5, GPT-5.4 mini, Gemini 3.5 Flash, Perplexity Sonar and four more.
SOURCES	17 of 72 answers, from four of twelve models. Most named: blog.salesflare.com, techradar.com.	17 of 72 answers, from four of twelve models. Most named: salesforce.com, blog.salesflare.com.	18 of 72 answers, from four of twelve models. Most named: ventureharbour.com, larksuite.com.

Negative shares are of each product’s own labels in that segment. A high share on a product with few labels is a warning; a low share on a product with many is salience, not sentiment. The rewording row names the models whose first choice held on the paraphrase.

Summary

The CRM category was asked as "CRM" across 12 models, six framings and three buyer segments. HubSpot CRM leads Small business at 50 percent of first choices and Mid-market at 41 percent. Salesforce Sales Cloud leads Enterprise at 58 percent. Each verdict is a clear leader. The leader changes once across the segments. Warnings concentrate on fit rather than function: complexity for small buyers, missing depth for large ones.

By buyer

Small business shows HubSpot CRM as endorsed leader, with Salesforce Sales Cloud warned against by 10 models at 76 percent negative and Dynamics 365 Sales at 82 percent. Mid-market places Salesforce Sales Cloud as a criticized challenger at 35 percent negative, with HubSpot CRM itself warned by 8 models. Enterprise reverses the pattern: Pipedrive at 75 percent negative, Zoho CRM at 50 percent, HubSpot CRM at 37 percent.

What to watch

First choices held under rewording for 4 of 12 models in Small business, 5 in Mid-market and 8 in Enterprise. The flips sit where nothing about the products changed, only the wording of the question. Sourcing stays thin: 17 of 72 answers in Small business and Mid-market, 18 of 72 in Enterprise, 4 models citing anything. Next edition should confirm whether the Enterprise stability persists.

Written from the figures on this page and checked against them; 0 sentences removed for stating a figure not in them.

SMALL BUSINESS · THE STANDING

ASKED OF THE STANDARD TWELVE MODELS

HubSpot CRM leads at 50%

Of 46 first choices across the direct, paraphrase, budget and scale prompts, over twelve models, asked on behalf of a small B2B company.

PRODUCT		SHARE		NEGATIVE LABELS	
■ ENDORSED LEADER		Above the 40% line, with a negative rate under a quarter.			
01	 HubSpot CRM		50%	12%	60
■ ACCEPTED CHALLENGER		Named, below the leader, and not argued against at scale.			
02	 Pipedrive		22%	6%	54
03	 Zoho CRM		11%	15%	55
04	 Salesflare		9%	0%	20
05	 Begin by Zoho		4%	0%	13
06	 Less Annoying CRM		2%	5%	19
07	 Freshsales		0%	3%	36
08	 EngageBay		0%	0%	10
□ CRITICIZED CHALLENGER		More than a quarter of their labels in this segment were negative.			
09	 Microsoft Dynamics 365 Sales		0%	82%	11
10	 Salesforce Sales Cloud		0%	76%	34

Share is the count of first choices over all twelve models. Every product with at least ten labels is listed; the rest are in standing.csv. The quadrant is the band, so it is stated once per band rather than once per row.

MOVEMENT

This is the first edition, so no change can be computed. From the next edition this block carries the move against the noise floor.

SMALL BUSINESS · WHAT THEY WARNED ABOUT

FOUR PRODUCTS, TWELVE MODELS

The leader at another segment is argued against here

A high negative share on a product with few labels is a warning; a low share on a product with many is salience, not sentiment. Each quote is the judge’s evidence for one label, with the model and the framing it came from.

 Salesforce Sales Cloud	26 of 34 labels negative · 10 of 12 models	76%
“Despite being marketed as "small business friendly," it's consistently described as a "time-sucking nightmare" for SMBs”		
MISTRAL SMALL · NEGATIVE PROMPT		
 Microsoft Dynamics 365 Sales	9 of 11 labels negative · 9 of 12 models	82%
“Like Salesforce, Dynamics is incredibly powerful but highly rigid and complex.”		
GEMINI 3.5 FLASH · NEGATIVE PROMPT		
 HubSpot CRM	7 of 60 labels negative · 7 of 12 models	12%
“flag enterprise-tier HubSpot as something small businesses should avoid if they do not need large-company features”		
PERPLEXITY SONAR · NEGATIVE PROMPT		
 Zoho CRM	8 of 55 labels negative · 7 of 12 models	15%
“reviewers sometimes mention pricing complexity, hidden fees, support quality, and setup/customization effort—so they’re worth a careful trial”		
GPT-5.4 MINI · NEGATIVE PROMPT		

The four products with the most negative labels in this segment. Every label and its evidence is in labels.csv, filtered to this category and segment.

Each quote is the judge’s evidence for one label, verbatim, with the model and the framing it came from. A quote without an attribution does not render.

SMALL BUSINESS · WHAT THEY CITE

FOUR OF TWELVE MODELS RETURNED SOURCES

17 of 72 answers came back with a source list

Citations where the model returns them, or the search results it consulted. Only a model with a search tool can return a source list, so these counts are drawn from four of the twelve and are not a share of all 72 answers.

SITES NAMED IN THE MOST ANSWERS	ANSWERS	PAGES	ANSWERS
blog.salesflare.com	6	blog.salesflare.com	5
techradar.com	6	/best-small-business-crm	
nutshell.com	6	crm.org	5
pcmag.com	5	/crmland/best-crm-for-small-business	
crm.org	5	pcmag.com	5
konabayev.com	5	/picks/the-best-small-business-crm-software	
cloudtalk.io	4	techradar.com	5
forbes.com	4	/best/best-crm-for-small-business	
		cloudtalk.io	4
		/blog/best-crm-software-for-small-businesses	
		forbes.com	4
		/advisor/business/software/best-crm-small-business	

Bars are answers, 0 to 17. A site cited twice in one answer counts once here; the cite count is in responses.csv.

The domain is set on its own line and the path below it, so a long URL breaks at a sensible point instead of mid-word.

WHICH MODELS CITE

The source lists came from Claude Haiku 4.5, Perplexity Sonar, Gemini 3.5 Flash and GPT-5.4 mini. The other eight models answered without returning one, so a site count here is a property of the tools, not of the field.

SMALL BUSINESS · HOW THEY ANSWERED

TWELVE MODELS, SIX FRAMINGS

Models one to five

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
A Claude Haiku 4.5	Salesforce Sales Cloud	HubSpot CRM, Pipedrive +2 Salesforce Sales Cloud		HubSpot CRM, Zoho CRM Wave		Microsoft Dynamics 365 Sales +1
GPT-5.4 mini	HubSpot CRM +3 Salesforce Starter / Sales Cloud	Pipedrive +2	HubSpot CRM +3	HubSpot CRM +1		LocaliQ +1
Gemini 3.5 Flash	Pipedrive +5	HubSpot CRM +3	Pipedrive +4	HubSpot CRM +3 Salesforce Sales Cloud +1	Pipedrive +5 Salesforce Sales Cloud	HubSpot CRM +6
Perplexity Sonar	Salesflare +5	Pipedrive +3	HubSpot CRM +5	HubSpot CRM +3		HubSpot CRM +2
Grok 4.1 Fast	HubSpot CRM +4 Salesforce Sales Cloud	HubSpot CRM +4 Salesforce Sales Cloud	HubSpot CRM +5	HubSpot CRM, Zoho CRM +4		HubSpot CRM +6

Four of twelve models held their first choice under the paraphrase: Grok 4.1 Fast, Llama 4 Maverick, GLM 4.7 FlashX and MiniMax M2.5. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

Models seven to eleven

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
 Mistral Small	Pipedrive +3	Salesflare +3 Salesforce Sales Cloud	Salesflare +6	HubSpot CRM +4	Microsoft Dynamics 365 Sales +1	Close +6
 DeepSeek V4 Flash	Salesflare +4 Salesforce Sales Cloud	HubSpot CRM +2	HubSpot CRM +4 Salesforce Sales Cloud	Zoho CRM +5		HubSpot CRM +4
 Llama 4 Maverick	HubSpot CRM +4	HubSpot CRM +3		Begin by Zoho +1		Salesforce Sales Cloud
 Qwen 3.7 Flash	HubSpot CRM, Pipedrive +1	HubSpot CRM +3	HubSpot CRM +3 Salesforce Sales Cloud	HubSpot CRM +2	Close, Pipedrive +3 Salesforce Sales Cloud	Microsoft Dynamics 365 Sales +3
 Kimi K2	HubSpot CRM, Salesflare +3	Pipedrive +2 Salesforce Sales Cloud	HubSpot CRM +5 Keap	Zoho CRM +3	Pipedrive +3 Salesforce Sales Cloud	HubSpot CRM +6

Four of twelve models held their first choice under the paraphrase: Grok 4.1 Fast, Llama 4 Maverick, GLM 4.7 FlashX and MiniMax M2.5. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

Models thirteen to fourteen

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
GLM 4.7 FlashX	HubSpot CRM +2	HubSpot CRM +4	Salesflare +5 Pipedrive	HubSpot CRM, Zoho CRM +2		HubSpot Sales Hub +4
 MiniMax M2.5	HubSpot CRM +2	HubSpot CRM +4	HubSpot CRM, Pipedrive +3 Salesforce Sales Cloud	Bigin by Zoho, Less Annoying CRM +3	+6 Microsoft Dynamics 365 Sales +1	ActiveCampaign +5

Four of twelve models held their first choice under the paraphrase: Grok 4.1 Fast, Llama 4 Maverick, GLM 4.7 FlashX and MiniMax M2.5. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

MID-MARKET · THE STANDING

ASKED OF THE STANDARD TWELVE MODELS

HubSpot CRM leads at 41%

Of 51 first choices across the direct, paraphrase, budget and scale prompts, over twelve models, asked on behalf of a mid-market B2B company.

PRODUCT		SHARE		NEGATIVE LABELS	
■ ENDORSED LEADER Above the 40% line, with a negative rate under a quarter.					
01	 HubSpot CRM		41%	19%	54
■ ACCEPTED CHALLENGER Named, below the leader, and not argued against at scale.					
02	 HubSpot Sales Hub		18%	0%	11
03	 Zoho CRM		10%	13%	54
05	 Pipedrive		2%	9%	46
06	 Freshsales		2%	10%	30
07	 Less Annoying CRM		2%	0%	11
08	 Nutshell		2%	0%	11
09	 Microsoft Dynamics 365 Sales		0%	7%	29
10	 Close		0%	8%	13
11	 Bitrix24		0%	20%	10
□ CRITICIZED CHALLENGER More than a quarter of their labels in this segment were negative.					
04	 Salesforce Sales Cloud		10%	35%	49

Share is the count of first choices over all twelve models. Every product with at least ten labels is listed; the rest are in standing.csv. The quadrant is the band, so it is stated once per band rather than once per row.

MOVEMENT

This is the first edition, so no change can be computed. From the next edition this block carries the move against the noise floor.

MID-MARKET · WHAT THEY WARNED ABOUT

FOUR PRODUCTS, TWELVE MODELS

The leader at another segment is argued against here

A high negative share on a product with few labels is a warning; a low share on a product with many is salience, not sentiment. Each quote is the judge’s evidence for one label, with the model and the framing it came from.

 **Salesforce Sales Cloud**

17 of 49 labels negative · 10 of 12 models

35%

“Don't buy Salesforce unless you're an enterprise with admin expertise and budget for the true TCO.”

DEEPSEEK V4 FLASH · NEGATIVE PROMPT

 **HubSpot CRM**

10 of 54 labels negative · 8 of 12 models

19%

“CRMs to Approach with High Caution... HubSpot: Tops lists with 300+ negative G2 reviews (capped), Risk Score 100/100”

GROK 4.1 FAST · NEGATIVE PROMPT

 **Zoho CRM**

7 of 54 labels negative · 7 of 12 models

13%

“Zoho CRM: 300+ negative G2 reviews... Affordable but frustrating for growing teams.”

GROK 4.1 FAST · NEGATIVE PROMPT

 **Pipedrive**

4 of 46 labels negative · 4 of 12 models

9%

“Avoid if you have marketing + sales + customer success teams... Pipedrive (7/100)”

MINIMAX M2.5 · NEGATIVE PROMPT

The four products with the most negative labels in this segment. Every label and its evidence is in labels.csv, filtered to this category and segment.

Each quote is the judge’s evidence for one label, verbatim, with the model and the framing it came from. A quote without an attribution does not render.

MID-MARKET · WHAT THEY CITE

FOUR OF TWELVE MODELS RETURNED SOURCES

17 of 72 answers came back with a source list

Citations where the model returns them, or the search results it consulted. Only a model with a search tool can return a source list, so these counts are drawn from four of the twelve and are not a share of all 72 answers.

SITES NAMED IN THE MOST ANSWERS	ANSWERS	PAGES	ANSWERS
salesforce.com	4	pcmag.com	3
blog.salesflare.com	4	/picks/the-best-crm-software	
monday.com	4	resources.rework.com	3
hubspot.com	4	/tools/crm/best-crm-software-2026	
pcmag.com	4	superoffice.com	3
nutshell.com	3	/blog/best-sales-crm	
clickup.com	3	blog.hubspot.com	2
resources.rework.com	3	/sales/crm-examples	
		blog.salesflare.com	2
		/best-b2b-crm	
		business.com	2
		/categories/crm-software	

Bars are answers, 0 to 17. A site cited twice in one answer counts once here; the cite count is in responses.csv.

The domain is set on its own line and the path below it, so a long URL breaks at a sensible point instead of mid-word.

WHICH MODELS CITE

The source lists came from Claude Haiku 4.5, Perplexity Sonar, Gemini 3.5 Flash and GPT-5.4 mini. The other eight models answered without returning one, so a site count here is a property of the tools, not of the field.

MID-MARKET · HOW THEY ANSWERED

TWELVE MODELS, SIX FRAMINGS

Models one to five

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
A Claude Haiku 4.5		Salesforce Sales Cloud		HubSpot CRM, Zoho CRM		
GPT-5.4 mini	HubSpot Sales Hub, Salesforce Sales Cloud +1	HubSpot CRM +2 Salesforce Sales Cloud	HubSpot CRM, Salesforce Sales Cloud +3	HubSpot CRM +4		
Gemini 3.5 Flash	HubSpot CRM +3	HubSpot Sales Hub +3	Salesforce Sales Cloud +4	Begin by Zoho, Zoho CRM +4 Odoo +2	HubSpot CRM, Salesforce Sales Cloud +1	Act! +10
Perplexity Sonar	HubSpot Sales Hub +4	HubSpot CRM +4		HubSpot CRM +3		
Grok 4.1 Fast	HubSpot CRM, Pipedrive +3 Salesforce Sales Cloud	HubSpot CRM +3 Salesforce Sales Cloud	+4	Zoho CRM +3 Pipedrive	HubSpot CRM, Salesforce Sales Cloud +2	Bitrix24 +6

Five of twelve models held their first choice under the paraphrase: Claude Haiku 4.5, DeepSeek V4 Flash, Llama 4 Maverick, Kimi K2 and GLM 4.7 FlashX. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

Models seven to twelve

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
 Mistral Small	HubSpot Sales Hub +3 Salesforce Sales Cloud	Salesflare +5	Salesforce Sales Cloud +9	Zoho CRM +3 HubSpot CRM		HubSpot CRM +2
 DeepSeek V4 Flash	HubSpot CRM +4 Salesforce Sales Cloud	HubSpot CRM +3 Salesforce Sales Cloud	+8	HubSpot CRM +4	Dynamics 365, HubSpot CRM, Salesforce Sales Cloud +1	Bitrix24 +5
 Llama 4 Maverick	Coevera +3	Coevera		Nutshell +4		
 Qwen 3.7 Flash	HubSpot Sales Hub +7	HubSpot CRM, HubSpot Sales Hub +5	+6	HubSpot CRM +3 AgileCRM		Act! +6
 Kimi K2	HubSpot Sales Hub +4 Salesforce Sales Cloud	HubSpot Sales Hub +3	+5	Freshsales, Zoho CRM +1 HubSpot CRM	HubSpot CRM, Salesforce Pro Suite +1	HubSpot CRM +4
GLM 4.7 FlashX	HubSpot CRM +7	HubSpot CRM +4	+8	HubSpot CRM +4		HubSpot CRM +2

Five of twelve models held their first choice under the paraphrase: Claude Haiku 4.5, DeepSeek V4 Flash, Llama 4 Maverick, Kimi K2 and GLM 4.7 FlashX. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

Model thirteen

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
 MiniMax M2.5	HubSpot Sales Hub +3	HubSpot CRM, Salesforce Sales Cloud		Begin by Zoho, Less Annoying CRM +3		 Close +8

Five of twelve models held their first choice under the paraphrase: Claude Haiku 4.5, DeepSeek V4 Flash, Llama 4 Maverick, Kimi K2 and GLM 4.7 FlashX. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

ENTERPRISE · THE STANDING

ASKED OF THE STANDARD TWELVE MODELS

The leader changes: Salesforce Sales Cloud at 58%

Of 43 first choices across the direct, paraphrase, budget and scale prompts, over twelve models, asked on behalf of an enterprise B2B company.

PRODUCT		SHARE		NEGATIVE LABELS	
■ ENDORSED LEADER		Above the 40% line, with a negative rate under a quarter.			
01	 Salesforce Sales Cloud		58%	19%	62
■ ACCEPTED CHALLENGER		Named, below the leader, and not argued against at scale.			
02	 Microsoft Dynamics 365 Sales		19%	6%	54
03	 Creatio		5%	6%	17
04	 HubSpot Sales Hub		5%	22%	18
06	 SAP Sales Cloud		0%	20%	25
07	 Oracle Sales		0%	20%	15
09	 SugarCRM		0%	20%	10
□ CRITICIZED CHALLENGER		More than a quarter of their labels in this segment were negative.			
05	 Zoho CRM		5%	50%	32
08	 HubSpot CRM		0%	37%	38
10	 Pipedrive		0%	75%	16

Share is the count of first choices over all twelve models. Every product with at least ten labels is listed; the rest are in standing.csv. The quadrant is the band, so it is stated once per band rather than once per row.

MOVEMENT

This is the first edition, so no change can be computed. From the next edition this block carries the move against the noise floor.

ENTERPRISE · WHAT THEY WARNED ABOUT

FOUR PRODUCTS, TWELVE MODELS

The leader at another segment is argued against here

A high negative share on a product with few labels is a warning; a low share on a product with many is salience, not sentiment. Each quote is the judge’s evidence for one label, with the model and the framing it came from.

 Pipedrive	12 of 16 labels negative · 9 of 12 models	75%
“Better suited for SMB and mid-market \u2014 generally not recommended for complex enterprise deployments.” DEEPSEEK V4 FLASH · DIRECT PROMPT		
 Salesforce Sales Cloud	12 of 62 labels negative · 9 of 12 models	19%
“Avoid Salesforce unless you need enterprise-grade customization that justifies unpredictable scaling costs” QWEN 3.7 FLASH · BUDGET PROMPT		
 HubSpot CRM	14 of 38 labels negative · 9 of 12 models	37%
“Avoid or be cautious with SMB/mid-market-focused CRMs... Lacks depth in advanced sales ops, territory management, and compliance for 100+ users” GROK 4.1 FAST · NEGATIVE PROMPT		
 Zoho CRM	16 of 32 labels negative · 8 of 12 models	50%
“performance degrades at 50+ users or millions of records... thinner compliance... create risks” GROK 4.1 FAST · NEGATIVE PROMPT		

The four products with the most negative labels in this segment. Every label and its evidence is in labels.csv, filtered to this category and segment.

Each quote is the judge’s evidence for one label, verbatim, with the model and the framing it came from. A quote without an attribution does not render.

ENTERPRISE · WHAT THEY CITE

FOUR OF TWELVE MODELS RETURNED SOURCES

18 of 72 answers came back with a source list

Citations where the model returns them, or the search results it consulted. Only a model with a search tool can return a source list, so these counts are drawn from four of the twelve and are not a share of all 72 answers.

SITES NAMED IN THE MOST ANSWERS	ANSWERS	PAGES	ANSWERS
ventureharbour.com	5	ventureharbour.com	5
larksuite.com	5	/best-enterprise-crm-software-solutions	
nutshell.com	5	articsledge.com	3
salesforce.com	4	/post/enterprise-crm-software	
g2.com	4	bitrix24.com	3
articsledge.com	3	/articles/top-15-enterprise-crm-software-in-2021.php	
bitrix24.com	3	clarity-ventures.com	3
clarity-ventures.com	3	/ecommerce-integrations/crm-enterprise-solutions-software	
		clarify.ai	2
		/blog/best-crms-for-b2b-sales-teams	
		clay.com	2
		/guides/best/crm-b2b-sales	

Bars are answers, 0 to 18. A site cited twice in one answer counts once here; the cite count is in responses.csv.

The domain is set on its own line and the path below it, so a long URL breaks at a sensible point instead of mid-word.

WHICH MODELS CITE

The source lists came from Claude Haiku 4.5, Perplexity Sonar, Gemini 3.5 Flash and GPT-5.4 mini. The other eight models answered without returning one, so a site count here is a property of the tools, not of the field.

ENTERPRISE · HOW THEY ANSWERED

TWELVE MODELS, SIX FRAMINGS

Models one to five

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
A Claude Haiku 4.5	Salesforce Sales Cloud +4	Salesforce Sales Cloud +3	+3	+3		Oracle's on Demand CRM +2
GPT-5.4 mini	Salesforce Sales Cloud +2 HubSpot CRM	Salesforce Sales Cloud +2	Salesforce Sales/Service Cloud +7	Microsoft Dynamics 365 Sales +1 HubSpot CRM		
Gemini 3.5 Flash	Salesforce Sales Cloud +3 Oracle Advertising and CX +1	Salesforce Sales Cloud +2 Oracle CX +1	Salesforce Sales Cloud +5	Creatio +3 HubSpot CRM +2	Salesforce Sales Cloud +2 Oracle CX +1	Airtable +12
Perplexity Sonar	Salesforce Sales Cloud +1 HubSpot Sales Hub +1	Salesforce Sales Cloud +1 HubSpot Sales Hub	Salesforce Sales Cloud +6	Microsoft Dynamics 365 Sales +1 Salesforce Sales Cloud		Bitrix24 +8
Grok 4.1 Fast	Salesforce Sales Cloud +2 Oracle CX +3	Salesforce Sales Cloud +2 HubSpot CRM +1	Salesforce Sales Cloud +4 Creatio +4	Microsoft Dynamics 365 Sales +3 Bitrix24 +3	Salesforce Sales Cloud +2	Bitrix24 +3

Eight of twelve models held their first choice under the paraphrase: Claude Haiku 4.5, GPT-5.4 mini, Gemini 3.5 Flash, Perplexity Sonar, Grok 4.1 Fast, Qwen 3.7 Flash, GLM 4.7 FlashX and MiniMax M2.5. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

Models seven to ten

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
 Mistral Small	Salesforce Sales Cloud +4	Microsoft Dynamics 365 Customer Engagement +3	Salesforce Sales Cloud +5 Pipedrive	Conduyt +2 HubSpot Sales Hub +1		HubSpot CRM +3
 DeepSeek V4 Flash	Microsoft Dynamics 365 Sales, Salesforce Sales Cloud +1 Freshsales +3	Salesforce Sales Cloud +1 HubSpot Sales Hub +2	Salesforce Sales Cloud +5	Microsoft Dynamics 365 Sales +1 HubSpot CRM +1	Microsoft Dynamics 365 Sales, Salesforce Sales Cloud HubSpot CRM +1	HubSpot CRM +7
 Llama 4 Maverick	Salesforce Sales Cloud, Zoho CRM +2	HubSpot Sales Hub, Salesforce Sales Cloud		Creatio +3	HubSpot Sales Hub +3	
 Qwen 3.7 Flash	Salesforce Sales Cloud +2	Salesforce Sales Cloud +3	Salesforce Sales Cloud +2 Zoho CRM	Zoho CRM +2 Salesforce Sales Cloud		Copper +5

Eight of twelve models held their first choice under the paraphrase: Claude Haiku 4.5, GPT-5.4 mini, Gemini 3.5 Flash, Perplexity Sonar, Grok 4.1 Fast, Qwen 3.7 Flash, GLM 4.7 FlashX and MiniMax M2.5. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

Models thirteen to fifteen

The first choice sits on the top line of each cell, the count of other products it named in gray below, and what it argued against in red.

MODEL	DIRECT	PARAPHRASE	COMPARATIVE	BUDGET	SCALE	NEGATIVE
 Kimi K2	Salesforce Sales Cloud +2	Salesforce Sales Cloud Enterprise/Un limited +1	Salesforce Sales Cloud +3  SAP Sales Cloud	Microsoft Dynamics 365 Sales +4  Salesforce Sales Cloud		 Close +4
GLM 4.7 FlashX	Salesforce Sales Cloud +2  Oracle CX +1	Salesforce Sales Cloud +1	Salesforce Sales Cloud +4	GoHighLevel +1		 Capsule CRM +6
 MiniMax M2.5	Salesforce Sales Cloud +1  Close +1	Salesforce Sales Cloud +3	Salesforce Sales Cloud +5	Microsoft Dynamics 365 Sales +2  Conduyt +2		

Eight of twelve models held their first choice under the paraphrase: Claude Haiku 4.5, GPT-5.4 mini, Gemini 3.5 Flash, Perplexity Sonar, Grok 4.1 Fast, Qwen 3.7 Flash, GLM 4.7 FlashX and MiniMax M2.5. The negative framing returns no recommendation, so that column holds only what the answer argued against. Every cell is one row of responses.csv joined to labels.csv.

METHOD, IN BRIEF

THE FULL METHOD IS AT [GTM-AI-INDEX.COM/METHODOLOGY](https://gtm-ai-index.com/methodology)

How this report was made

The run

Every model in the tier was asked the same six framings of the buying question for each buyer segment, in a fresh session with web search on. No model was told about the index, and no answer was retried.

The judge

One model, Claude Opus 5, labels every product named in every answer against the same rubric, and attaches the sentence it was named in. Share counts first choices only, across the direct, paraphrase, budget and scale prompts.

The vendor table

A versioned map from the names a judge wrote to canonical products, with category-scoped readings and exclusions. This report was scored under v2026-09-15.2, and a category is a clear leader when the top share exceeds 40%, consensus when every model made the same product its sole first choice on the direct prompt, and contested otherwise.

The framings

Direct, paraphrase, comparative, budget-constrained, scale-constrained and negative. The paraphrase is the same question reworded; the negative asks what to avoid and returns no recommendation.

The noise floor

A calibration repeat asks the same prompts again with nothing changed. On the latest repeat, 66% of first choices changed; the floor for a category leader's share is 10 points. No move between editions is called a move unless it clears that floor.

What this is not

Nothing here is a recommendation by the index. It is a measurement of what these models said when asked, in this edition, under this vendor table.

THE TWO REPORTS

Standard Edition Pack

The twelve standard models, at all three buyer segments. Its standing is the public index's standing, computed the same way, so every figure can be checked against the site.

Expanded Edition Pack

The same twelve plus six flagship models. Each tier is computed on its own models; a combined eighteen-model standing appears only where the report says so, and never on the public site.

Every folio and cross-reference is derived from one page map, so the extent follows the content: one page per section, and the how-they-answered matrix paginated at six models a page. This report runs to 22 pages. Type is set in CSS pixels at ninety-six to the inch, so a 12px table cell prints at 9pt and nothing in the report is smaller.

CITE THIS EDITION

GTM AI Recommendation Index, September 2026 Edition, CRM. gtm-ai-index.com. Scored under vendor table v2026-09-15.2. Published under CC BY 4.0.

THE FILES BEHIND IT

standing.csv, labels.csv, responses.csv, prompts.csv, models.csv and index.sqlite, with the data dictionary in README.md.