medianet.

PHASE 3 · PUBLIC SECTOR EDITION

Opening the Black Box.

A framework for measuring reputation in LLM environments.

PHASE 3: An analysis of visibility, sentiment, and citation sources across Australian public sector programs in major AI models.



amec AWARDS | 2026
GLOBAL COMMUNICATION EFFECTIVENESS

SILVER WINNER · 2026 AMEC GLOBAL AWARDS
Best Measurement of the Impact of Communication on AI Answers

Executive summary.

When we released our first Black Box report at the end of 2025, we set out to apply traditional media-analysis rigour to a technology that had quietly become a major reputational channel.

We approached **LLM responses like any other media channel** : monitor it, code it, compare it, ask better questions of it.

Phases 1 and 2 showed how LLMs treat brands, while Phase 3 examines what happens when LLMs become the first point of contact for **government programs with high public and media scrutiny..**

We found that LLMs mirror the tone and sentiment of the prompt, so negatively framed questions get negative answers. We also continued to see a lag in uptake of new information such as new programs and policies. And we also found that, although government domains lead citations (48%), commercial providers supply more than a quarter (27%), the second-largest group.

This requires public sector teams **to monitor platforms individually**, engage cited providers, and publish clear, updated explainers on official sites immediately upon reform announcements. Then measure what the models say after each change, and keep iterating.

> *For essential government programs, AI answers are the new front counter.*
>
> *This poses a risk of citizens making decisions based on outdated or commercially filtered information, but it is also a powerful opportunity to shape the script and provide clear, timely and trusted guidance to the public.*



Amrita Sidhu
Managing Director · Medianet

Recap from Phase 1 and 2

LLMs favour owned content and soften criticism.

SOURCES

67%

of citations came from owned content, led by government sources at 44%

LLMs go straight to the source.

Traditional media supplied just 16% of citations. 45% of CommBank mentions repeated messaging from its own website word for word.

SENTIMENT

97%

of Westpac mentions in LLM answers were favourable, against 60% in mainstream media

LLMs are more skew positive.

Models weighed pros and cons and softened controversy. Vic Gov drew unfavourable coverage in over 50% of media reports but only 14% of LLM mentions.

Recap from Phase 1 and 2

AI visibility follows its own rules.

SHARE OF VOICE

26%

share of model for Youi in car insurance, against 7% of media coverage

Share of voice ≠ Share of Model.

NRMA led the media with 31% of coverage but held 20% share of model. Strong media performance doesn't guarantee a place in the answer.

LIFESPAN

Years

how long issues stay in LLM answers, against 2–4 weeks in media

LLMs have a very long memory.

Royal commissions, ASIC rulings and outdated web pages still shaped how models described organisations, long after the news cycle moved on.

Phase 3

Scope and methodology.

During Phase 2 of our research we noticed that government sources made up 44% of all citations, even though Government was not the subject of the prompts. Phase 3 therefore, puts the public sector at the centre with a focus on mentions of NDIS and My Aged Care.

These large federal programs were selected because they are in the middle of major reform and face heavy public scrutiny. We tracked mentions on LLMs for both across 30 prompts on ChatGPT, Google AI Overviews and Perplexity four weeks between September and October 2026. This produced 2,610 responses and 14,661 citations.

Our prompts were highly curated and designed to emulate real-world user behaviour and surface how LLMs represent the NDIS and My Aged Care. We also deepened our analysis to look at how the tone and wording of each prompt shaped the answer, specifically incorporating "sceptical" (negatively framed prompts) and "supportive" (positively framed prompts) into our framework.

LLM RESPONSES ANALYSED

2,610

1,305 per sector, across ChatGPT, Google AI Overviews and Perplexity.

CITATIONS CATEGORISED

14,661

Coded by source type, sentiment and program visibility.

ANNOUNCED · 2026

NDIS.

A new set of reforms was announced in 2026, including the Thriving Kids program. Still at the policy stage.

IMPLEMENTED · NOV 2025

My Aged Care.

Reforms took effect in November 2025, replacing systems many older Australians already knew.

Phase 3

Key findings.

Phases 1 and 2 showed how LLMs frame brands. Phase 3 asks a harder question. What happens when the subject is a government program that people depend on?

When someone asks an AI assistant how to get disability support or aged care, the answer is often their first stop. It tells them where to go, what to expect and what the program stands for. Government agencies write only part of that answer.

> *The risk is not an LLM criticising the program. The risk lies in LLMs providing inaccurate information to the public, and which sources the model uses to respond.*

01 **LLMs mirror the tone of the prompt.**
When prompts were framed negatively or sceptically, LLMs reflected the same tone in the answer instead of neutralising it.

02 **The lag cuts both ways.**
New government policy was missing from responses. So was new criticism of it.

03 **Government sources lead, followed by businesses and providers.**
Providers supplied 27% of citations, even on questions about what a program is.

04 **The platform changes the output.**
The more sources used by a model in an answer, the more diluted the official narrative becomes.

01

Sentiment.

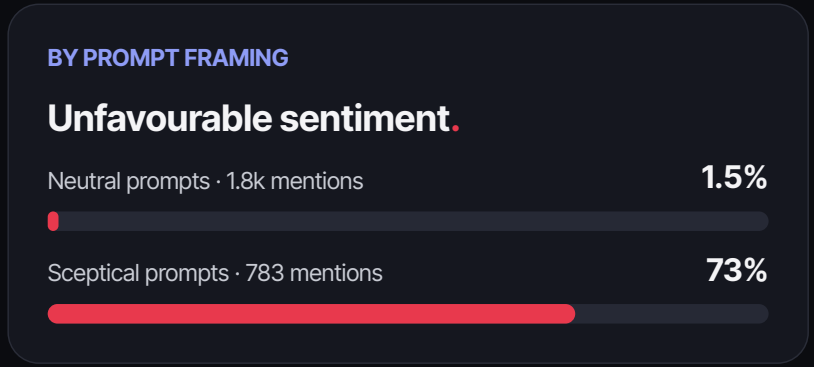
The sentiment of the prompt shaped the sentiment of the response.

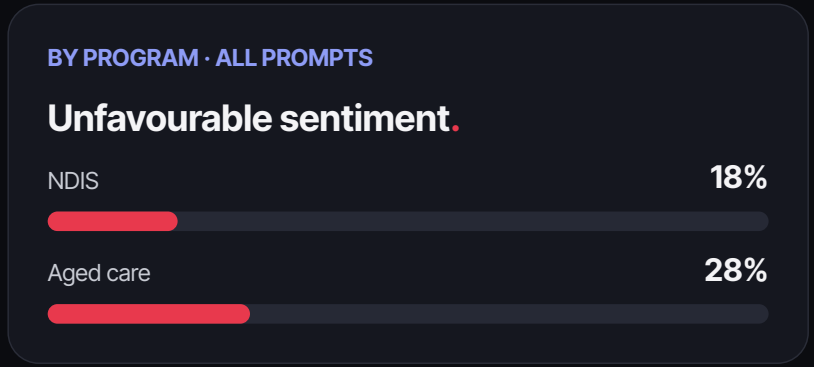
Overall, LLMs were favourable or neutral in 82% of NDIS responses and 72% of My Aged Care responses, which aligns with our previous findings around LLMs neutralising contentious issues.

For this phase, however, we took a closer look at the link between sentiment and the type of prompt being entered into an LLM. These results told a more nuanced story.

From prompts that were neutrally framed, LLMs almost never raised an issue. We asked 'How does the government treat people living with disability in Australia?' Not one NDIS reference was unfavourable. But when a prompt was framed negatively or sceptically, LLMs did not balance the response. Instead, the response reinforced and reflected the sentiment in the prompt. Just above 80% of NDIS and 69% of My Aged Care mentions were unfavourable in sentiment when responding to a negatively framed question or prompt.

This asks communication professionals to think more broadly about what influences their brand's sentiment beyond the input feeding the responses: who is entering these prompts, and are they more likely to ask sceptical questions?





02

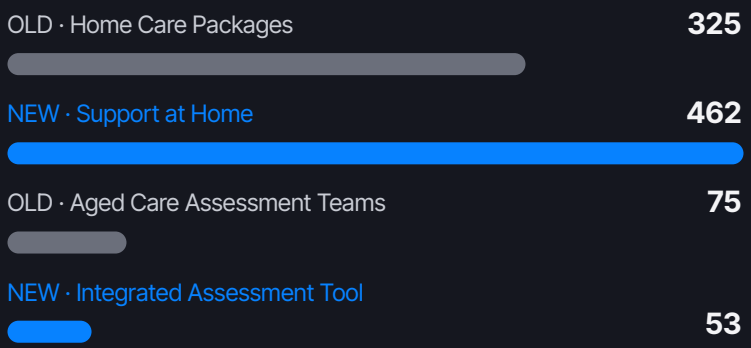
The traditional media lag.

Changes are slow to be picked up by LLMs.

Our research to date showed LLMs are slow to pick up emerging media themes, and slow to forget old ones. We interrogated this further by looking at how the media "lag" plays out in public sector narratives. The aged care department and NDIS have both undertaken legislative reform and transitioned to new programs, but each are on a different timeline. The aged care department has already implemented its new systems, while NDIS changes are newer and still at the policy stage.

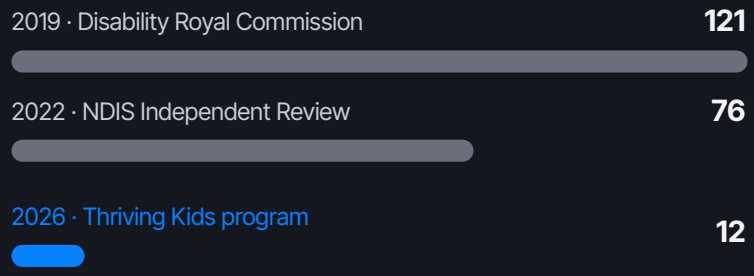
My Aged Care: The new system has arrived, but the old one hasn't quite left yet.

Australia's aged care system saw a raft of reforms implemented in November 2025 under new 'Support at Home' legislation. This system has replaced the previous "Home Care Packages" system, and while it was more prominent in responses, the old program was still being referenced.



NDIS: New policy overshadowed by older news.

Major service reforms to the NDIS, including a new "Thriving Kids" support program, were announced in April 2026. Unlike aged care services, old and new programs didn't sit side-by-side. Instead, LLMs characterised the scheme through older reference points. This included contentious issues such as the 2019 Disability Royal Commission and 2022 NDIS Independent Review. Changes announced in 2026 barely registered despite being an ongoing theme of media scrutiny.



Program mentions across responses

The risks of the lag.

MY AGED CARE

325

responses still mentioned Home Care Packages, against 462 for Support at Home

The risk: Communities engaging with aged care services may be confused by the old and new services contained in the same response.

NDIS

12

responses mentioned Thriving Kids, against 121 for the 2019 Disability Royal Commission

The risk: Information seekers will receive information about the current state of the scheme, without exposure to important service delivery changes.

11 months

A new name can take up to a year to be embedded in LLM models.

Almost 11 months later, Support at Home has only now started to emerge in LLM outputs. NDIS can prepare by updating official materials now, not after implementation.

Information accuracy management is as important as sentiment and reputation management.

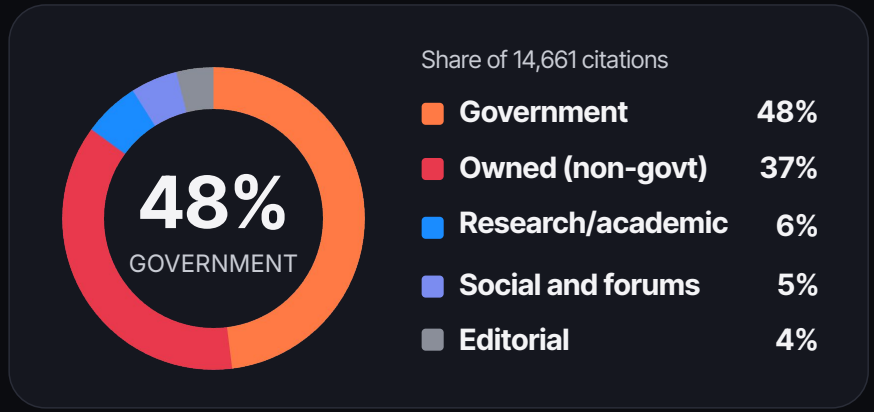
When something changes, provide explainers that name the old system, name what is replacing it, and give the date it changed.

03

Sources.

LLMs go straight to the source. For public programs, there are many.

Owned digital content from government domains predominated in citations. Of these URLs, 29% came straight from the NDIS or My Aged Care websites, reinforcing the influence of information directly from the relevant departments.



Owning the domain: A double-edged sword for public sector

The sheer volume of government digital resources can be a blessing and a curse. Government domains have a larger share of citations, but they also have a vast network of complex digital content, with many stakeholders with different priorities contributing to this ecosystem. The Department of Health's URL was the leading unique domain to be cited, meaning a high portion of the resources generating LLM outputs had been prepared by a different department's team with their own communications strategy.

This creates more complexity around programs and services that are already difficult to communicate. Public-facing programs and services need to be delivered in **simple and accessible** content, which is challenging to manage amid the sheer volume of government content being surfaced.

The type of prompt influences citations

Phase 3 confirmed the predominance of owned content among LLM citations. However, there are exceptions to this rule.

Prompts: Sceptical prompts get more non-government answers. When prompts asked about problems, concerns, challenges or barriers, only 31% of citations came from the government in comparison to 52% for neutrally-framed questions.

Sources: Editorial citations appeared more in responses that contained criticism. Only 23% of overall responses were unfavourable, however 63% of editorial responses carried an unfavourable sentiment. Phase 3 research showed that **when an LLM is asked a negatively framed question, they are more likely to reach for journalism.**

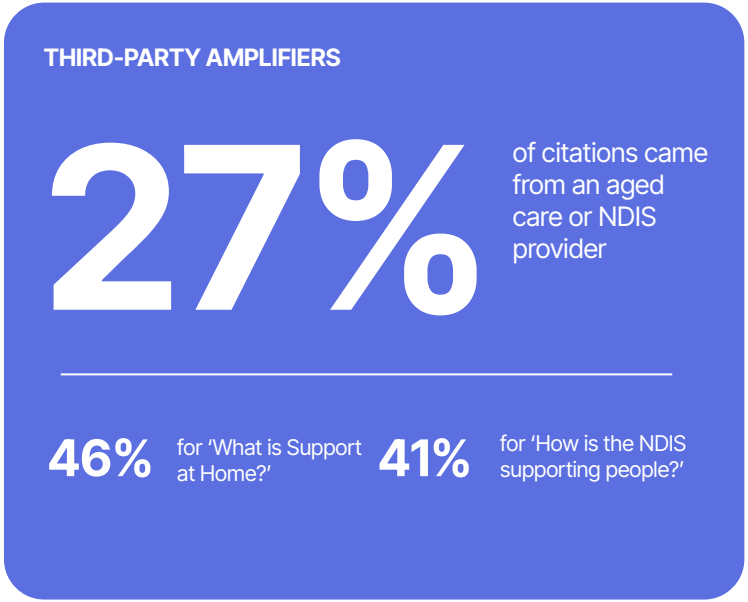
The power of third-party sources.

For a corporate brand, owned content mostly means your own website. For a government program, a whole commercial layer sits between the agency and the public.

Providers are just as likely to publish clear, crawlable explainers used by LLMs, but government agencies do not control what those businesses say about eligibility, fees, or reforms.

NDIS and My Aged Care providers accounted for **27% of overall citations**—making them the **second leading group driving citations**—and were often cited in prompts that weren't specifically asking for advice.

Providers supplied 46% of citations in response to "What is Support at Home?" and 41% for "How is the NDIS supporting people?", indicating this stakeholder group is vital in explaining public programs.



04

Models.

Different platforms will change the response.

To date, Medianet has reviewed LLM outputs as a whole ecosystem. We took a more granular approach in Phase 3 to identify any correlations or departures between our three models - ChatGPT, Google AI Overviews, and Perplexity. Our results demonstrate the need to interrogate LLM Visibility across different platforms, rather than one blended number.

3 citations per response

Government **88%**

Owned (non-govt) **7%**

ChatGPT averaged three citations per response, with the vast majority coming from the public sector. Just 5% came from providers, indicating the government voice gets the most cut-through on this model.

6.8 Citations per response

Government **36%**

Owned (non-govt) **46%**

Almost half of citations in Google AI Overviews derived from provider content. This model also drew from a wider range of sources, with an average of 6.8 citations referenced per response. Google AI Overviews helped amplify the profile of aged care services and tended to surface references to older programs the most.

7 citations per response

Government **43%**

Owned (non-govt) **42%**

Perplexity had the highest number of citations per response (7 citations) and produced the highest number of responses with an unfavourable tone. This model also tended to reference editorial content slightly more than the other models, possibly contributing to the higher proportion of mixed sentiment.

New programs can appear well represented in one platform, but poorly in another. A wider number of citations can potentially dilute government sources as the more places an LLM looks for information, the less frequently it will repeat the official line.

Key learnings

What we learned in Phase 3.

01 Accuracy can be a reputation risk.

For a corporate brand, an outdated LLM answer can cause reputation damage. For a government service, it can send a citizen to the wrong place for information. Monitoring LLM answers should be considered by service delivery teams, not just those managing the comms.

02 Announcements are not enough.

NDIS changes were announced in April, and have barely been picked up by LLMs 6 months later. Despite a high volume of traditional media activity around the changes, these did not translate into LLMs. New reforms need clear and dated explainers on the agency's own pages, where the LLMs are likely to go first.

03 Stakeholders are harder to control.

Service providers represented a high number of citations. These helped communicate information around NDIS and aged care programs, but will often do this in their own words. Engage with those being most frequently cited in LLM outputs to ensure their messaging and information aligns with your organisation.

Recommendations

What this means for public sector communication.

Phases 1 and 2 demonstrated how LLM Visibility can shape reputation and brand perceptions. Phase 3 has shown that LLMs are as important when communicating new government services and programs. There is, however, a wide range of stakeholders competing for visibility, and a large portion of these (providers) sit outside the government's immediate control.

Check what LLMs tell the public.

Run questions people are likely to ask about a government service, for example how to apply, w*hat it costs, how to complain*. Check answers for accuracy, not just sentiment.

Measure by platform and message.

Track each platform separately and how each talks about your key services, including the counter-narratives that dilute the official line.

Publish explainers from the moment something changes.

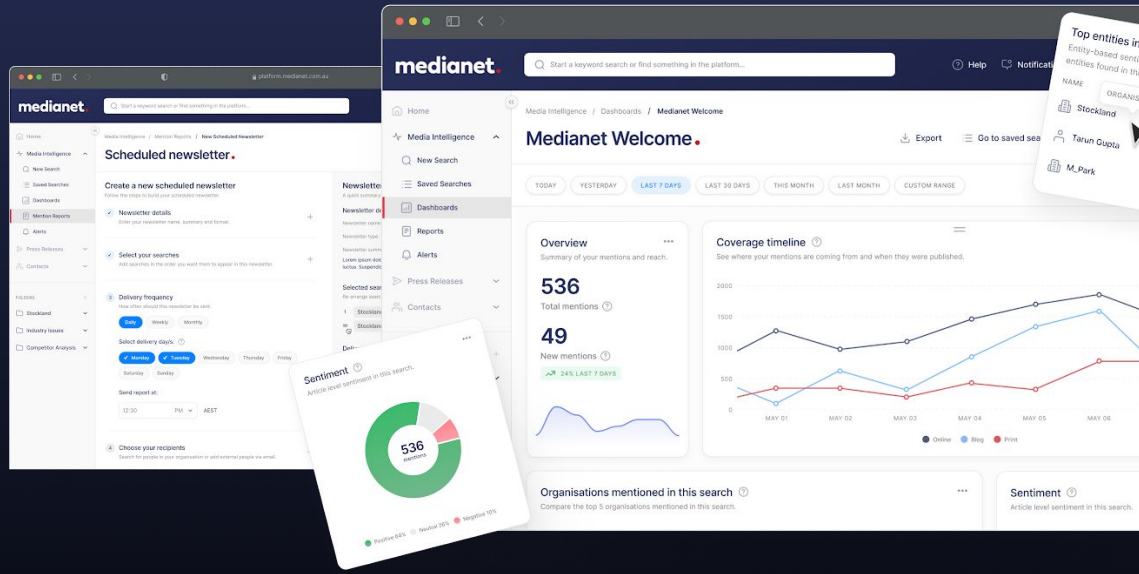
For each reform, publish a dated, plain-English page on the program's own site on or before announcement day. Remove or clearly label older versions.

Plan for the lag.

Assume LLMs describe your program as it was months ago. Brief frontline and contact centre staff to expect citizens quoting outdated AI answers.

About Medianet

Medianet is Australia's most trusted media intelligence solution that helps service delivery, whole of government and regulatory organisations make better communication decisions and understand the public perception of their policies and operations. Learn more here.



medianet.

1300 616 813
news@medianet.com.au
www.medianet.com.au

amec Member

medianet.

Have you joined Leadership in the Loop?

A dedicated series of intimate networking events for ***communication leaders.***

The Leadership in the Loop event series, hosted by Medianet, is a space to discuss reputation at the intersection of technology, communication and leadership. Join the waitlist to be invited to our exclusive **invite-only events** with high-calibre guests, panelists and other communication leaders.

Join the waitlist

