KIT

The Case for Support Architecture

What this is	Fund development as a capability rather than a writing task: the case components written once and adapted per funder, an evidence register where every claim carries its source, a go / no-go standard applied before writing begins, the attachment set kept current, and the handoff that puts delivery staff in the room before submission.
It assumes	You seek grant or contract funding. You do not need a development director — the smaller the shop, the more of this is currently in one person's head.
Change before use	Every component text and every claim. The architecture transfers; your case does not come from anyone else.

THE ARCHITECTURE IS THE SPREADSHEET

This document explains what a case for support has to answer and how the parts are maintained. **The architecture itself is the accompanying spreadsheet** — the component library with an owner and a last-reviewed date, the evidence register linking every claim to its source, the go / no-go scorecard, the prospect register, the attachment currency sheet, and the submission calendar.

The scorecard computes a recommendation and the summary counts the claims with no source and the attachments out of date. Read this once. Work in the spreadsheet.

The problem this fixes

Fund development gets treated as a writing task because the writing is the visible part. But a proposal is a claim about what an organization can do, and the claim is only as good as the design behind it and the capacity beneath it. Organizations buy the writing, win sometimes, and end up with an award to deliver and no more capability than they started with.

The symptom is specific and recognizable: every proposal is written from scratch against the funder's guidelines, because there is no case the organization holds in common. The person who writes well becomes the bottleneck, then the single point of failure, and when they leave, the pipeline leaves with them. Nothing accumulated.

> *A good writer cannot invent a program in an organization that has no capacity to implement it. Where that happens, the award is won and the organization is worse off.*

The seven questions a case has to answer

Every funder asks these in their own words and their own order. Answer them once, well, in components you own — then adapt, rather than rewrite.

1. **Whose problem, stated as they experience it.** Not the sector's problem in the abstract. The specific difficulty of specific people, described in a way they would recognize.
2. **Why it persists.** The mechanism, not the lament. A problem with no stated mechanism invites the reader to supply their own, which is usually that nobody has tried hard enough.
3. **What you do about it.** The service model in enough detail to be operated: what happens, how often, for how long, delivered by whom.
4. **Why that should work.** The chain from what you do to what changes, and what has to be true for it to hold. This is the logic model, in prose.
5. **How you will know.** The indicators, their definitions, and an honest statement of what your data can and cannot demonstrate.
6. **Why you.** Standing, track record, position, and relationships — stated as facts a reader could check, not as adjectives.
7. **What it costs, and what happens after.** The true cost, the request, the rest of the funding, and what continues when this award ends.

The seventh question is the one most often answered badly, and it is the one a program officer remembers. "We will seek continued funding" is not an answer. Naming the specific continuation path — or saying plainly that the work is time-limited and why that is appropriate — is.

Components, not documents

The case is maintained as components rather than as a finished proposal, because no funder ever wants the finished proposal you already have. Each component has an owner, a word budget, and a last-reviewed date.

Component	Typical budget	Reviewed
Organizational overview and standing	150–300 words	Annually
The problem, in the population's terms	200–400 words	Annually, or when the data changes
Why it persists — the mechanism	150–250 words	Annually
Service model description, per program	300–600 words	When the model changes
Theory of change in prose	150–300 words	When the logic model changes
Measurement and evidence approach	200–350 words	When indicators are revised
Capability, track record, and partnerships	200–400 words	Annually
Sustainability and what continues	150–300 words	Annually
Governance, leadership, and staffing	150–250 words	On any leadership change
Equity, access, and language approach	150–300 words	Annually

The discipline that makes this work is the review date, not the writing. A component library nobody refreshes becomes a library of statements that were true once, and those are worse than no library at all — because they get submitted.

The evidence register

Every factual claim in the case carries a source, an as-of date, and the person who verified it. This is the sheet that prevents the most damaging category of error, which is not exaggeration but decay: a number that was accurate when it was written, is now three years old, and is still being submitted.

- A claim with no source does not go in the case. Not "we will find it later" — out, until it has one.
- Every claim carries the date it was verified, and an expiry appropriate to the kind of claim it is. Population figures, service volumes, and staffing counts all decay at different rates.
- Claims about your own performance come from your own systems, and the sheet names the field and the report they come from, so anyone can re-derive the number.
- Claims about the sector or the population are cited to their source and reviewed on the same cycle. If it cannot be sourced, it is stated as your observation and marked as such.

The register makes a due-diligence request or a monitoring question a lookup rather than an archaeology project.

The go / no-go standard

Applied before anyone writes anything. The purpose is to make declining a considered act rather than an omission, and to stop the organization from pursuing what is open instead of what fits.

Criterion	The question	A no-go looks like
Programmatic fit	Does this fund work we would do anyway, at a scale we intend?	A new program invented to match a solicitation
Delivery capacity	Could we staff, supervise, document and evidence this, starting on their timeline?	The capacity check has a No with no date beside it
True cost against the award	Does the award cover the real cost, including the share of overhead it consumes?	The gap is real and there is no named source to cover it
Compliance load	What does this award oblige us to do, report, and retain, and can we?	Requirements we have no system for and no plan to build one
Relationship and standing	Do we have standing with this funder, or a credible path to it?	A cold application against a closed field
Opportunity cost	What does the staff time spent here not get spent on?	The submission displaces a renewal or a report that is due
Probability, honestly	What do we actually think the odds are, and on what basis?	Nobody will say a number out loud

The workbook scores these and computes a recommendation. The recommendation is advisory and the organization can overrule it — but the overrule is recorded with a reason, which means the pattern becomes visible over a year. An organization that overrules its own standard every time does not have a standard; it has a form.

The handoff that prevents the first-report surprise

The people who will deliver the work belong in the room **before** submission, not after the award. The workbook records four sign-offs against each pursuit, and none of them is ceremonial.

1. **Program** confirms the service model described is the one that will actually be delivered, at the volume stated.
2. **Finance** confirms the budget reflects true cost, that the match or cost-share is real and identified, and that the cash timing is survivable.
3. **Compliance** confirms the reporting, retention, and eligibility obligations are ones the organization can meet, and names what has to be built if not.
4. **Leadership** confirms the opportunity cost is acceptable and that the organization intends to do this work.

A pursuit that cannot collect these four before the deadline is telling you something about the deadline, or about the pursuit.

Where the line is

This is an architecture for building your own capability. It is not a substitute for one, and it will not write anything for you. Every component in it has to be written by someone who knows the organization, and the whole point of holding it as components with owners and review dates is that the knowledge stays with the organization when a person leaves.

It is also not advice on solicitation registration, charitable-solicitation filings, professional-fundraiser licensing, or the tax treatment of any gift or award. Those depend on your state and your circumstances and belong to your own counsel.

Published by The CARES Collaborative, a program of Knotts Family Services, Inc., for adaptation and use by charitable and community-serving organizations. No registration, no fee, no attribution required. This document is a working template, not legal, tax, or accounting advice — the requirements that apply to your organization are confirmed by your own counsel, auditor, and licensing authority. Version 1.0 · August 2026 · carescollaborative.org/open-shelf