

What the Community Needs to Decide

DAWN BoF Chairs

Wes Hardaker

Adrian Farrel

Recall What We are Doing Today

- The purpose of a WG-forming BoF is (loosely)
 - Is there a problem that needs to be solved?
 - Is it in scope for the IETF?
 - Is there a community that wants to work on it?
- And
 - Is there enough focus that a working group has a chance of delivering in a reasonable time-frame
 - Note that there is a lot of pressure in the industry to make progress around AI agents protocols
- The chairs can only give advice and guidance
 - The community will decide what it wants to work on
 - The IESG will decide whether the proposed work is reasonable
- The chairs (with guidance from the IESG) advise “focus”
 - The IESG is likely to reject a charter that is not focused on a small number of quick deliverables
 - Future topics can be added as “not currently in scope” but added in a later re-charter

Big Questions

- There are a few questions that need to be kept in mind when discussing the charter
 - How the community decides to answer these questions will shape the charter put to the IESG
- 1. Should we limit to discovery of “AI Agents”, or generalise applicability to similar “Entities”?
- 2. What are the first critical discovery items that will bootstrap future needs?
- 3. Should we limit to a “single organisation” or allow “Internet-wide” discovery?
- 4. Should we jump in with DNS as a starting point, or investigate other options before starting protocol work?
- ...

AI Agents or General Entities?

Entity: A system component that communicates with other entities in a peer-to-peer or client-server relationship. Entities include, but are not limited to, AI agents, tools, skills, tasks, compute workloads, software services, task owners, network functions, and application endpoints.

(<https://datatracker.ietf.org/doc/draft-farrel-dawn-terminology/>)

- Consider:
 - The pressing need is for discovery of AI agents
 - Building re-useable tools with wide applicability is a Good Idea
 - Engineering to enable generic solutions may result in over-engineering
 - Defining all of the possible entities could take a lot of effort
- Chairs' recommendation
 - Scoping in all types of entity may make the charter too wide to be acceptable
 - Make AI Agents the initial focus, but keep an eye open for general applicability

What are the first critical discovery items?

- Discovery ecosystem includes (at least):
 - Registration and management of an entity to be discovered
 - Search based on a formalised encoding of search criteria
 - Discovery lookup to discover an agent/entity via a minimum viable data set
 - Retrieval of additional information
 - Capability exchange and negotiation
- Consider:
 - There is a need for an architecture to cover all of this (is it IETF work?)
 - Each piece might need standardisation work (is it all IETF work?)
 - The discovery lookup is a separable functional component
 - We should aim for a bootstrap to get things up and running soon
- Chairs' recommendation
 - Scoping in all elements of the ecosystem would make the charter too big to be acceptable
 - Focus on what the first critical discovery items that will bootstrap future needs
 - Keep a list of future work items for a possible re-charter
 - Consider using other working groups for other topics (existing or new groups)

Single Organisation or Bigger Picture?

- Deployment model is debatable:
 - Domains of specific AI agent types
 - Private domains (such as enterprises / organisations)
 - AI agents located within a service provider or data centre
 - AI agents anywhere in the Internet
- Consider:
 - It is important to “walk before we run”
 - Scoping out some deployment models could result in failure to adopt
- Chairs’ recommendation
 - The deployment models listed above should be in scope
 - If specific problems or distinctions are discovered they should be called out

Start with DNS or Look at All Options?

- There are a **lot** of potential solutions
 - Partial, non-definitive list at:
 - <https://github.com/danielkinguk/discovery/blob/main/solutions-list.md>
 - See also:
 - <https://datatracker.ietf.org/doc/draft-jimenez-dawn-discovery-landscape/>
- Consider:
 - The IETF has a name service lookup protocol (DNS)
 - Re-using protocols is quickest way to deployment and benefits from proven implementations
 - Over-loading a protocol can make it bloated and even unstable
 - Inventing new protocols is a lot of fun!
- Chairs' recommendation
 - Make initial attempt to use DNS (plus extensions) as the discovery bootstrap
 - Pay attention to short-comings and failings
 - Be informed by all other work

Recommended Approach

- The chairs recommend (straw man)...
 - As much focus as possible
 - Small list of priority deliverables
 - List items that are “out of scope for now”
 - But allow them to be brought into scope through re-chartering in the future
- Don’t forget that other urgent work could be the subject of a parallel charter effort (leading to one or more further WGs)
- The community must decide and “negotiate” with the IESG
 - Discuss the charter in this meeting and on the mailing list