

FIA Investigators' meeting

Colorado State University - Lory Student Center
Cherokee Ballroom – Upper Level room 236D
451 Isotope Drive, Fort Collins, CO 80523

April 19/20 2012

Meeting Format

To help refine the focus of this meeting, we sent out a list of 7 discussion questions, which are included at the end of this agenda. The goal is to use these questions to shape the meeting.

- We will begin with a 30-minute presentation from each project, which will give the projects a chance to discuss the relevant aspects of their projects.
- Following that part of the agenda, we will have a group discussion of each question in turn, calling on the various projects as appropriate to fill in details and facilitate a “compare and contrast” exercise. We anticipate that this part of the meeting will occupy us for the first afternoon and the second morning.
- At the beginning of the second day, we have set time aside for each of the projects to give a further brief presentation, perhaps in response to the questions and discussion that has occurred on the first day. It is often the case that after an evening to think about what has been said, it becomes clear what sort of clarification and elaboration will be helpful.
- Following this discussion, we have set aside time for our VID participants and outside observers to give us reactions, and share any thoughts they have based on their own experience.

Agenda – Thursday, April 19

Day 1: Cherokee Ballroom 236D, upper level, Lory Student Center

8:00-9:00 Buffet breakfast available in Room 230

9:00-9:15 Introduction and summary of objectives

Part 1: project presentations

9:15-9:45 Named Data Networking

9:45- 10:15 XIA

10-15-10:45 MobilityFirst

10:45-11: 00 Break, in Room 230

11:00-11:30 Nebula

11:30-12:00 ChoiceNet

12:00-1:00 Lunch

1:00-1:30 Introduction of outsiders and VID participants

Part 2: Topic discussions

1:30-5:00 discussion of questions

2:30-3:00 Break in Room 230

5:45 Board bus for dinner outside the Lory Student Center

6:00 Board bus for dinner at Hilton Ft. Collins

6:30 Dinner at Fort Collins Country Club

Agenda – Friday, April 20

Day 2: Cherokee Ballroom 236D, upper level, Lory Student Center

8:00-9:00 Breakfast in Long's Peak Room 264

9:00-10:15 Day two reprise of project presentations: 10-15 minutes each.

- Named Data Networking
- XIA
- MobilityFirst
- Nebula
- ChoiceNet

10:15-10:30 Morning Break outside meeting room

10:30-12:00 Discussion continued

12:00-1:00 boxed lunch available in Long's Peak Room 264

1:00-1:15 Discussion of future FIA meetings (NSF)

1:15-3:00 Presentation/responses from VID participants/outsideers

3:00-3:30 time buffer allowance

3:30 adjourn

The timing of the discussions is hard to predict, as is the timing of the various presentations on day 2. So the overall timing may be variable, but we will work for adjournment by 3:30.

The discussion questions

Who are the service providers in your architecture, and what is the resulting provider ecosystem? (Some of the FIA architectures seem to presume a provider ecosystem similar to today: a connected set of packet forwarders. Some presume other services related to carriage, such as storage providers.)

- What is the incentive of each of these actors to enter into their line of business? Where would your architecture require payments among actors to sustain viability?

Options for control: which actors can influence the behavior of a transfer?

- Does your architecture provide user control over aspects of service selection: routes, service qualities, or providers of support service (e.g. like DNS in today's Internet)?
- To what extent does your architecture support or resist the goals of those who wish to control access to classes of information (e.g. governments, rights-holders). How does this position influence the balance of power in your network, and its viability? Which actors have the ability (or perhaps the *easy* ability) to block communication among willing end-points?
- IP addresses accidentally turned out to be scarce resources, for no good reason. What features of your architecture might turn out to be "scarce resources" or resources over which some potentially powerful actor could exercise control?
- Do you have hierarchies with single points of control at the root? Is there information you share with partners that has to be signed by a trusted third party?
- Are there policies that you have explicitly embedded in your design?

What is the range of services that the system provides to the higher layers?

- Compared to today's Internet, would you expect the same sort of commercial entities at the higher layers?
- For example, (especially in the context of those architectures that emphasize information retrieval), would you imagine that there would be CDNs operating on top of your architecture?
- Does your architecture provide an API that defines the service interface of your system?

Interfaces among providers

- What types of information is expected to be exchanged between providers? This goes beyond packet forwarding to include:
 - Routing information
 - Naming information (e.g. DNS zone transfers)
- An interconnection agreement between providers in today's Internet may have Service Level Requirements, or specify aspects of routing policies (could

- potato, hot potato). What would you expect to find in inter-provider agreements in your architecture?
- To what extent do services provided to higher levels (see above) require negotiation or cooperation among the various actors that make up the overall network?
 - What mechanisms does your architecture provide for negotiation among service providers?
 - What range-of-functions are supported by the protocols and mechanisms that hook them together?
 - Operators are sometimes worried about all getting together to solve operational issues. It is hard to do and looks like anti-trust. What are the "top five" aspects of your architecture that require operational coordination?

Market forces and regulation

- To what extent does your proposal facilitate or limit the use of competition as a discipline on the market?
- If regulation were proposed to require some sort of non-discriminatory access or "network neutrality", what might that mean in your design? Where might forms of discriminatory service emerge?

Evolvability

- How does your architecture allow innovation and the migration to new mechanisms?
- Which sorts of evolution seem to require global coordination, like the migration to IPv6 today?

Trust, isolation and availability

- What sorts of trust assumptions does your design make about the various actors that make up the ecosystem?
- Does your architecture provide means for instrumentation or data-gathering? What sorts of data? Internal structure of the network, usage, routes, outages, etc?
- To what extent does your architecture include tools to detect that actors are not functioning properly? Which actors have access to these tools?
- How do your options for control allow different actors to respond to actors that are not trustworthy or mis-functioning?
- Availability often implies "extra" or "diverse" resources. Does your architecture depend on resources that are otherwise under-utilized to achieve high-availability. Is economics a barrier to a high-availability network? Both within a region and across regions, does your design allow the operator to trade off explicitly between cost and availability/resilience?