

Infrastructure-to-application Information Exposure

I2AEX BoF

IETF 83

Monday, March 26, 2012 – Paris

I2AEX

Why?

- To set distributed applications free to get smarter

How?

- Enabling infrastructures to expose useful information

What?

- Extensions / Usage modes for existing protocols

But Why?

- You (!) keep asking that
 - In ALTO
 - CDN usage for ALTO, ALTO for high capacity networks, s2c notifications...
 - Software Driven Networking initiatives
 - BoF @ IETF-82, sdnp@ietf.org
 - Datacenter initiatives
 - dc@ietf.org, dcrg-interest@irtf.org, bar BoFs...
- Maybe we can **help (some) apps get smarter**
 - Focus restriction: fully and partially controlled env
 - CDN, datacenter apps


And How?

- Identifying what kind of **i2a** information would be useful
 - Focus restriction reminder: in fully-to-partially controlled environments
- Identifying the **smallest possible tool** useful to the **widest possible audience**
 - (If any)

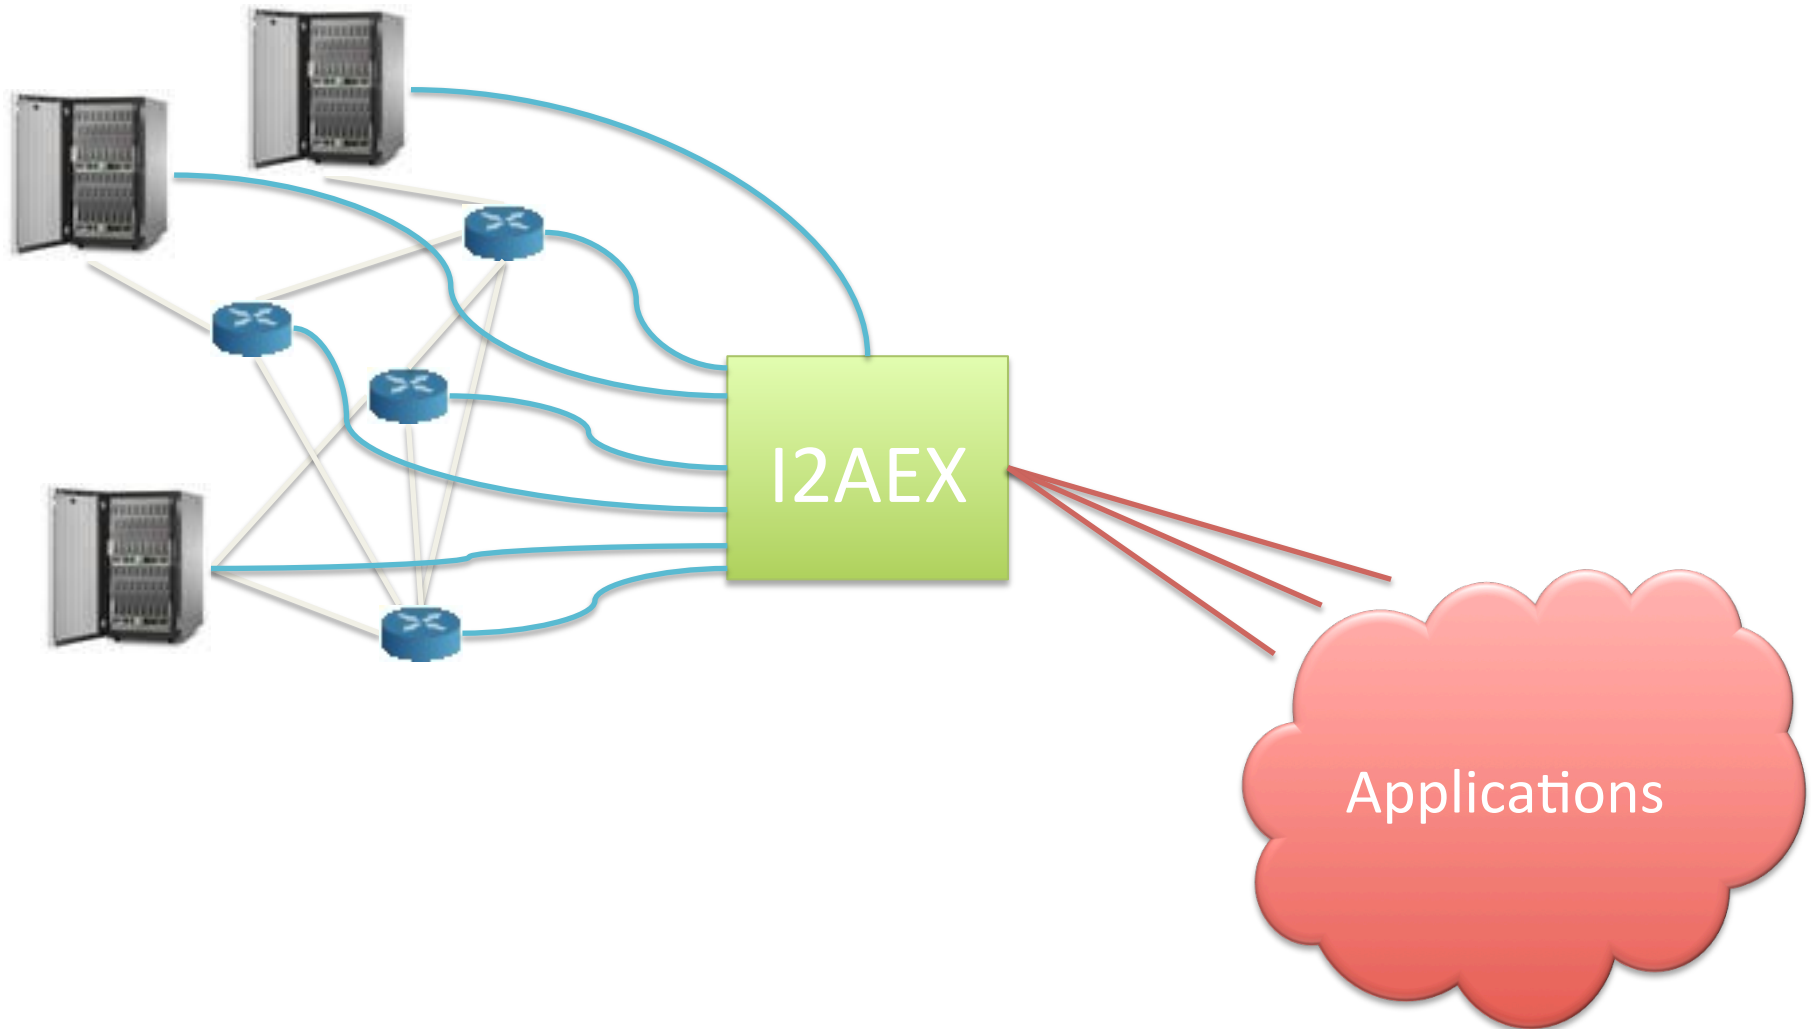
So What?

- Possible extensions / usage modes for existing protocols
 - Easily accessible from app layer
 - Reminder: focus on fully/partially controlled apps, **CDN** and **datacenters** to begin with
 - Info types:
 - Network load (on “reasonable” timescales)
 - Resource availability (CPU, storage)
 - Content availability
 - ...

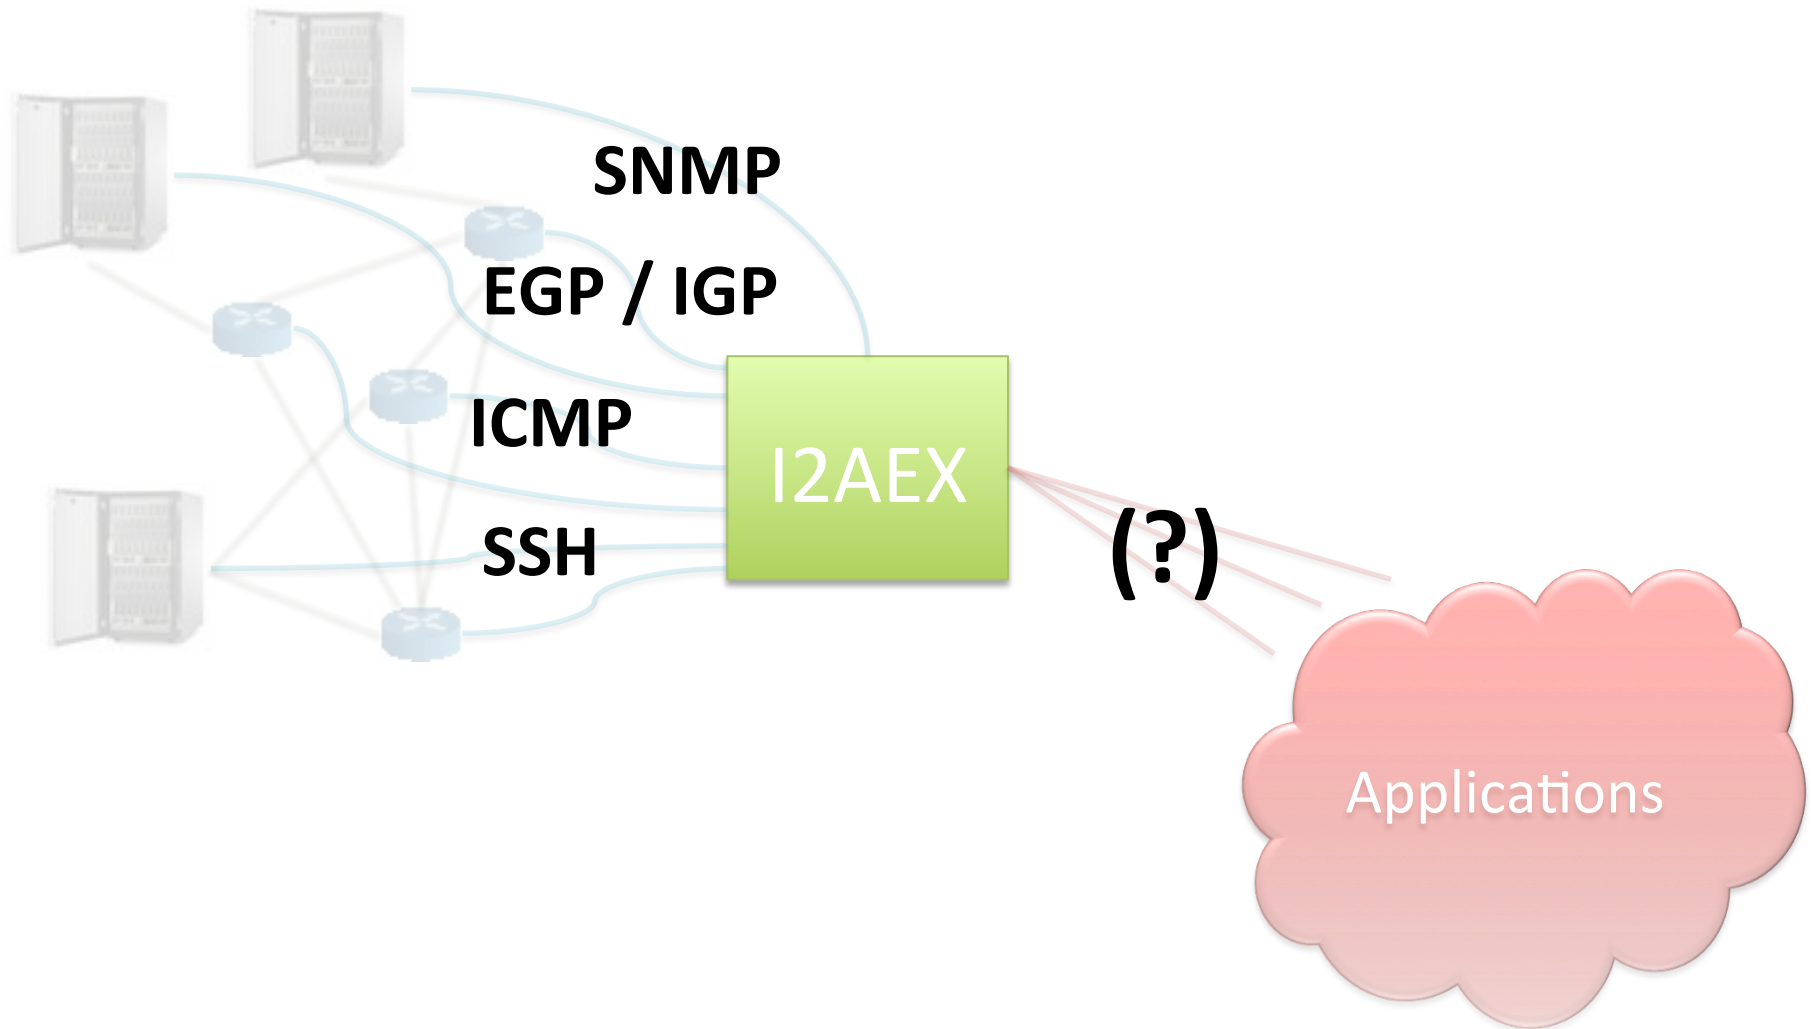
Starting Point



Goal (?)



Goal



First Question

Is there a need for (yet!) another infrastructure-to-application communications mechanism?

– [Focus restrictions reminder goes here]

Doing it the ALTO way

The ALTO Protocol is...

- Network-to-application Information Exposure
 - Very “application friendly” (HTTP+REST+JSON)
 - Intended to let network operators (and whoever else wants to) provide information about the more and less appropriate network destinations
 - E.g. “you may not want to get your file from there, you fool, that’s Australia!”
- Easily extensible
- Done (OK, almost...)

The ALTO Protocol is also...

- Primarily intended for uncontrolled apps
 - Untrusted parties / Low confidentiality
- Primarily designed for quasi-static info
 - Based on network and cost maps

ALTO i2aex Extensions

- New information types
 - Network load
 - Resource availability
 - ...
- Server-to-client notifications
- Incremental updates

Incremental updates

- draft-pbryan-json-patch (APPSAWG item)

ALTO Information

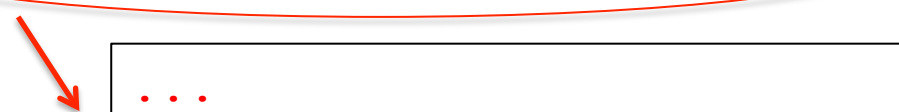
```
{
  "resources": [
    {
      "capabilities": {
        "cost-modes": [
          "numerical"
        ],
        "cost-types": [
          "routingcost"
        ]
      },
      "media-types": [
        "application/alto-costmap+json"
      ],
      "uri": "http://alto.ietf.org/costmap/num/routingcost"
    },
    ...
  ]
}
```

```
{
  "meta" : {},
  "data" : {
    "cost-mode" : "numerical",
    "cost-type" : "routingcost",
    "map-vtag" : "1266506139",
    "map" : {
      "mypid1": { "mypid1" : "0",
                  "peerpid2" : "10",
                  "pid3" : "20"},
      ...
    }
  }
}
```



New Information Types

```
...  
  "media-types": [  
    "application/alto-costmap+json",  
  ],  
  "uri": "http://alto.ietf.org/costmap/num/routingcost"  
},  
{  
  "media-types" : [  
    "application/alto-serverload+json"  
  ],  
  "uri" : "http://alto.tilab.com/serverload"  
},  
...
```



```
...  
  "server1": {  
    "lastmin" : "0.7",  
    "last5min" : "0.3",  
    "last15min" : "0.2"  
  },  
...  
}
```


Server-to-client Notifications

- HTTP polling?
- Websockets?
- ALTO-over-XYZP?

```
"media-types": [  
    "application/alto-costmap+json"  
],  
"uri": [  
    "http://alto.ietf.org/costmap/num/routingcost",  
    "xyzp://alto.ietf.org/costmap/num/routingcost"  
],
```

...

Second Question

IFF answer to First Question is yes, then

does the ALTO protocol provide a reasonable basis for an i2a information exposure mechanism? (Focus restrictions apply.)