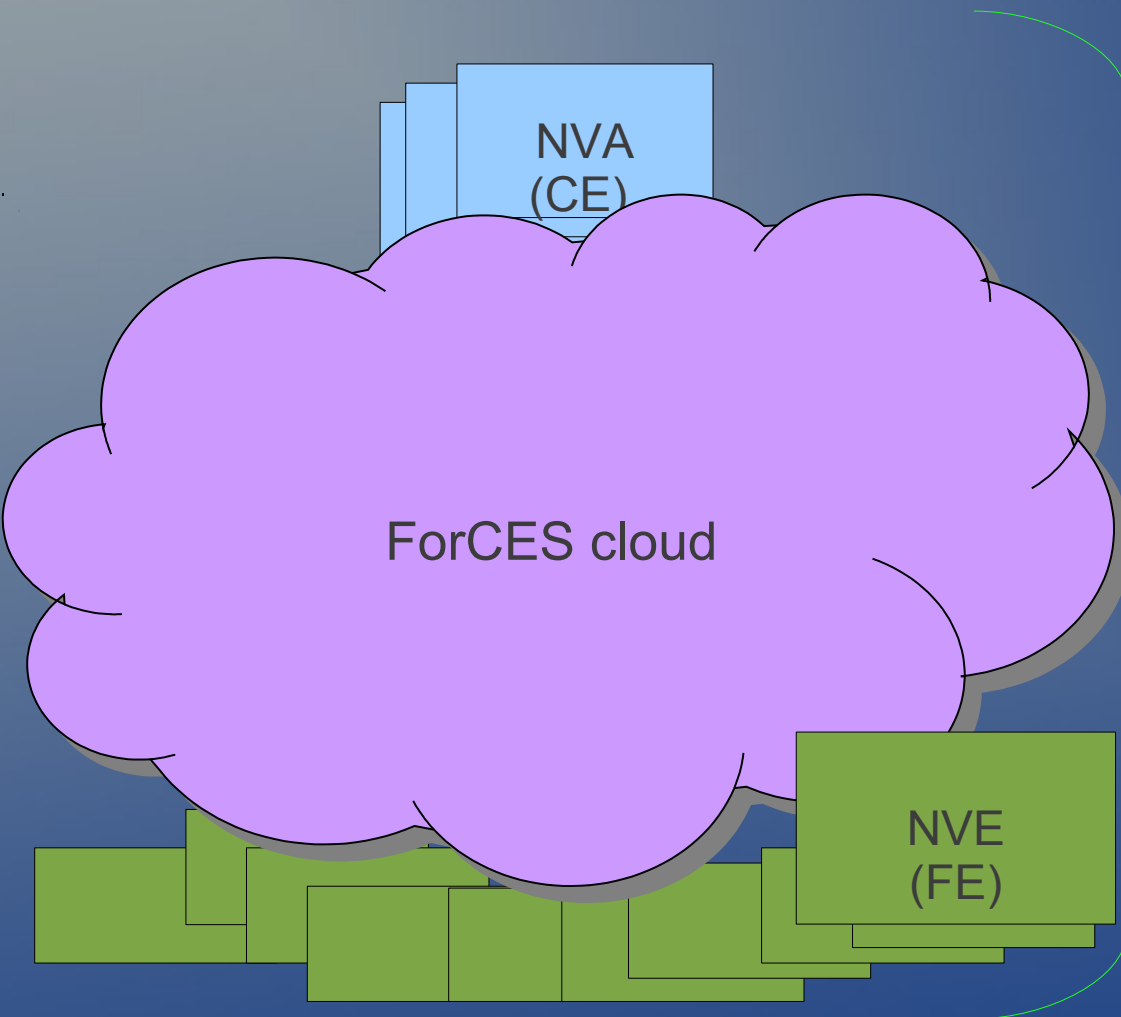


ForCES nvo3
IETF 88, Nov 4 2013
NVO3 WG
Jamal Hadi Salim
Bhumip Khasnabish

ForCES: Functional scope



- Network Element (NE)
 - Packet Processing Entity
 - Constitutes CEs and FEs
 - Multiple CEs to FE for HA
 - CE/FE Physical or virtual

- NE components distributed
 - Local within box/rack/room
 - Geographically distributed
 - Within rooms/buildings
 - Across the internet

ForCES Architecture In A Nutshell

- A binary protocol (The *Verbs*)
 - A modular transport for the protocol
- A data model (The *nouns*)
 - Logical Functional Block constructs
- Combine the above and you have a language
 - [*<verb> <noun> [args]*]⁺
 - Few verbs but infinite possibilities of nouns

Protocol Semantics

- Simple Verbs
- Transactional capability (2PC)
- Various Execution modes
- Scalability via batching and command pipeline
- Security
- Traffic Sensitive Heartbeating
- High Availability

Data Modeling

- Based on XML (RFC 5812)
- Respect for backward and forward compat
 - Let the CE deal with disparate versions
- Main constructs
 - *Datatype*: definitions used by LFBs
 - *LFB Classes*: Basic packet processing entity
 - *Components*: Control entities a CE is aware of
 - *Capabilities*: define LFB capabilities
 - *Events*: define events that an LFB can generate

LFB model Example

Datatype definition

```
<dataTypeDef>
  <name>foobartype</name>
  <synopsis>Describes The foobar</synopsis>
  <struct>
    <component componentID="1">
      <name>foo</name>
      <synopsis>that foo</synopsis>
      <typeRef>uint32</typeRef>
    </component>
    <component componentID="2">
      <name>bar</name>
      <synopsis>the bar</synopsis>
      <typeRef>uint32</typeRef>
    </component>
  </struct>
</dataTypeDef>
```

MyLFB

Components

```
<component componentID="1" access="read-write">
  <name>foobar</name>
  <synopsis>The Foo and Bar thingy</synopsis>
  <typeRef>foobartype</typeRef>
</component>
```

Capabilities


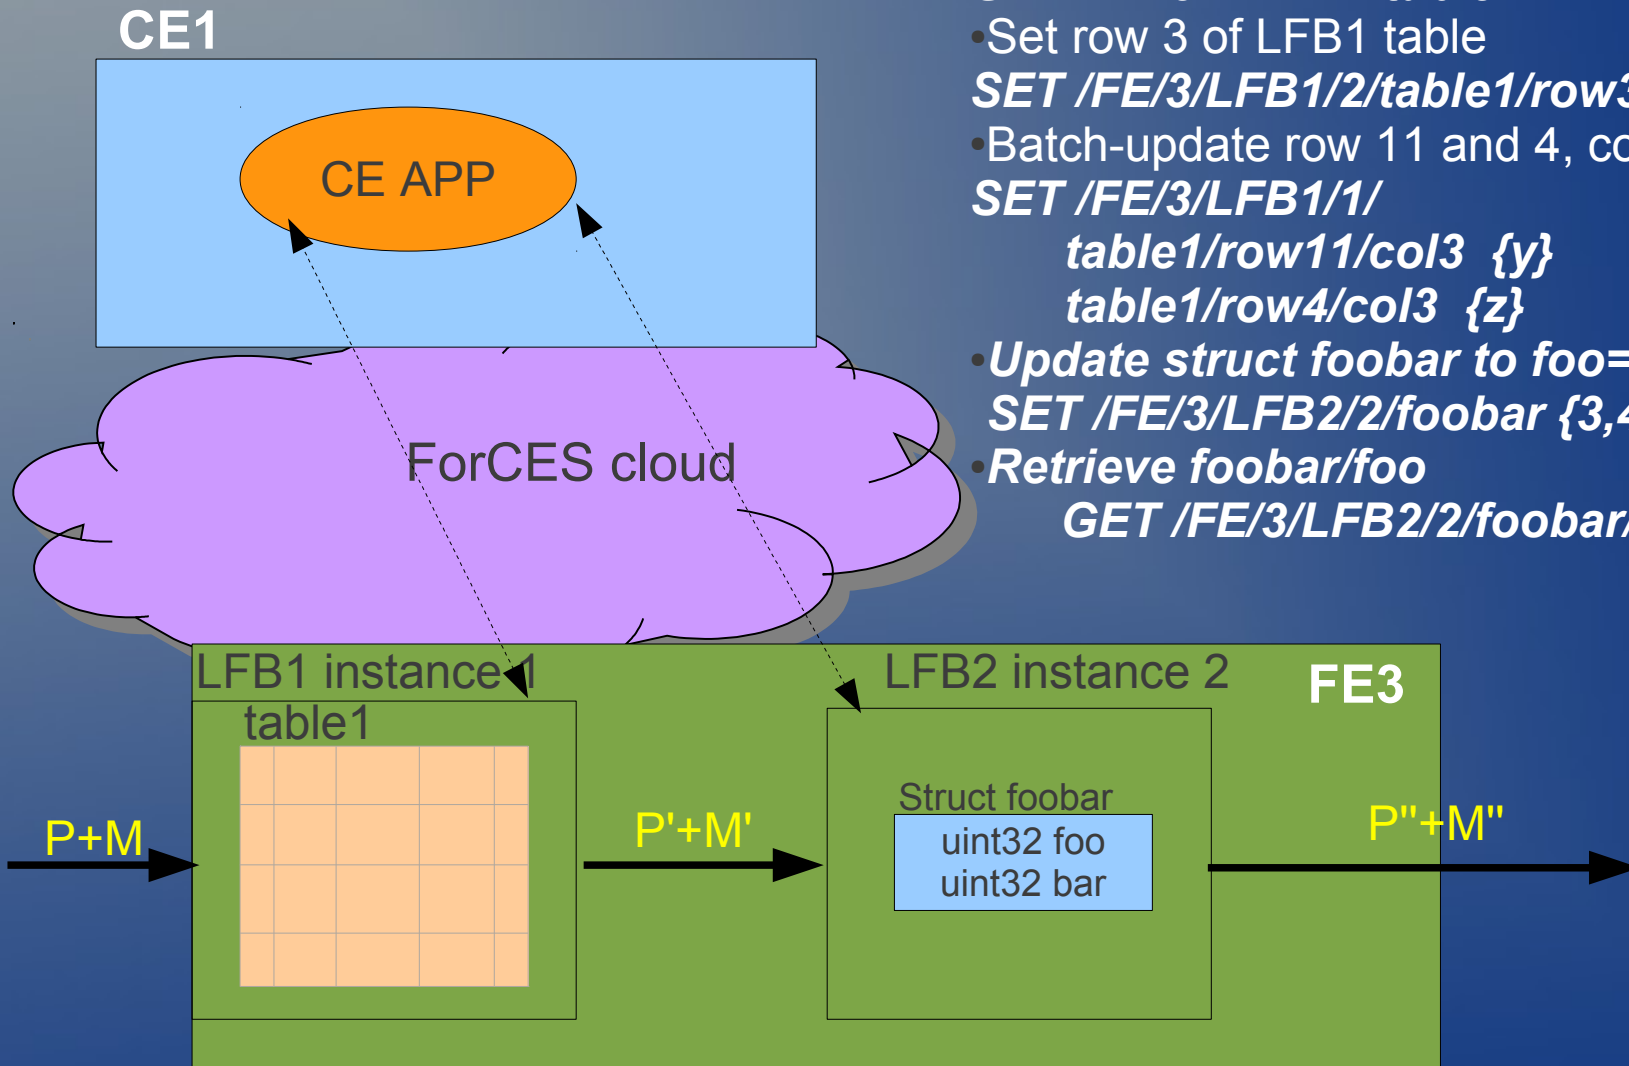
```
<component componentID="100" access="read-only">
  <name>foobarcount</name>
  <synopsis>The Foo and Bar capacity</synopsis>
  <typeRef>uint32</typeRef>
</component>
```

Events

```
<component eventID="1" access="read-only">
  <name>foobarwatch</name>
  .... watch foobar changes and report to ce.....
</component>
```

General Control Example

- Dump the LFB1 table
GET /FE/3/LFB1/1/table1
- Set row 3 of LFB1 table
SET /FE/3/LFB1/2/table1/row3 {a,b,c,d}
- Batch-update row 11 and 4, col 3 of the LFB1 table
**SET /FE/3/LFB1/1/
table1/row11/col3 {y}
table1/row4/col3 {z}**
- Update struct foobar to foo=3, bar=4
SET /FE/3/LFB2/2/foobar {3,4}
- Retrieve foobar/foo
GET /FE/3/LFB2/2/foobar/foo



Meeting nvo3 NVA/NVE requirements

- Less complex
 - Not a protocols menu (OF, netconf, ovsdb etc)
 - Extensible API based on model definition
 - CLI or control apps get a consistent interface
 - Agnostic whether we target over/underlay
 - Future revision require changes only in the model
 - Backward and forward compatibility built in
- Built in VNE High Availability support

Meeting nvo3 NVA/NVE requirements

- Efficient protocol (binary)
 - Designed to be fast and scale
 - thousands of NVEs
 - Millions of VNIs
 - Fast updates
 - Bulk table creates/updates/deletes/gets
 - In the 10-100K table updates per second possible
 - Fast acquisition of state and events distribution

Meeting nvo3 NVA/NVE requirements

- A data model (The *nouns*)
 - Can model variety of VNE types and constructs
 - Mapping tables
 - Scalar control knobs
 - Events vs redirects
 - Send control BPDUs to VNA
 - Send events or whole frames to VNA for policy decisions
 - Capacity and capability definitions
 - VNE type capabilities
 - Table and other capacity advertisement

Bits n Bites etc

- At this meeting there is a demo with Vms, virtual switches etc
 - Come to the ForCES meeting next session
 - Come to the bits and bites session thursday