

# The Software in SDN

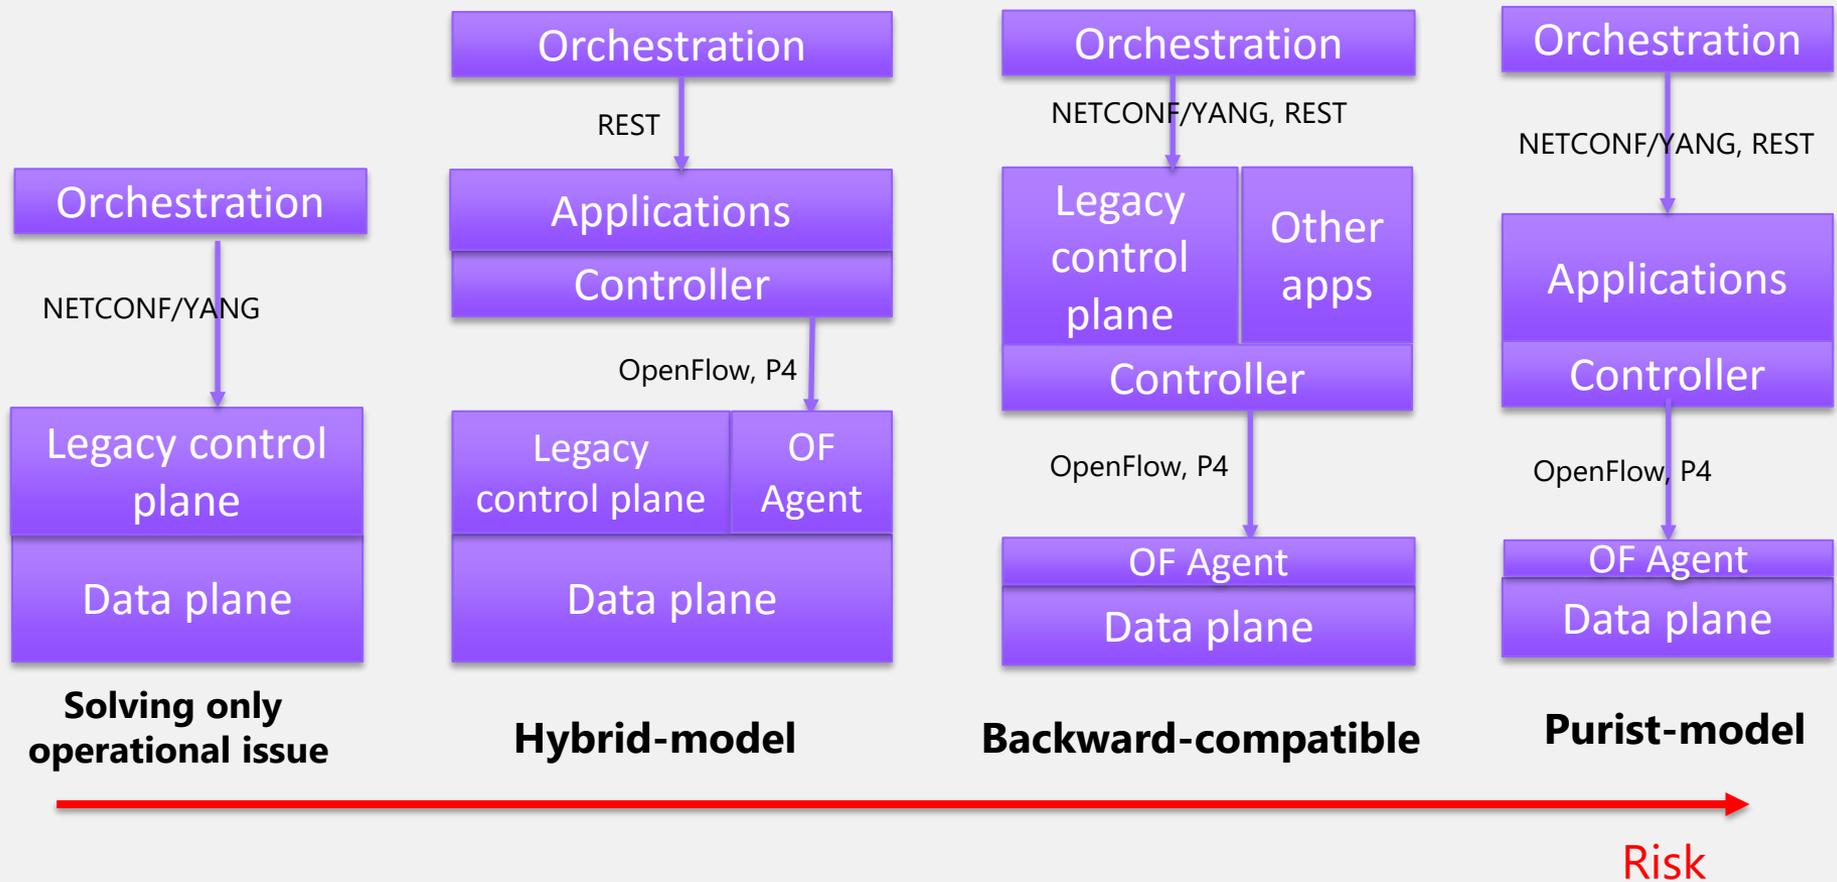
Programming in Opendaylight

# Presenter's Note

This presentation borrows heavily from a presentation on App Development by Srini Seetharaman and are available at

<http://sdnhub.org>

# OpenDaylight embraces the full spectrum of choices



# Lithium: List of Projects (Total: 40)

Project keyword	Description		
aaa	AAA Project	opflex*	OpFlex protocol
alto*	Application-Layer Traffic Optimization	ovsdb	OVSDB protocol and OpenStack integration
bgpcep	BGP/PCEP protocol library	packetcable	PacketCable PCMM/COPS
capwap*	Control and Provisioning of Wireless Access Point	pss*	Persistence Store Service
controller	OpenDaylight Controller	plugin2oc	Southbound plugin to the OpenContrail platform
defense4all	Radware Defense4All	reservation*	Dynamic Resource Reservation project
didm*	Device Identification and Drive Management	sdninterfaceapp	SDNi Cross-controller interface
discovery	Discovery Service	sfc	Service Function Chaining
dflux	OpenDaylight UI	snbi	Secure Network Bootstrap Infrastructure
docs	Documentation Project	snmp4sdn	SNMP4SDN Plugin
groupbasedpolicy	Group Based Policy Plugin	snmp*	SNMP Southbound Plugin
integration	Integration Framework	sxp*	Source-Group Tag eXchange Protocol
iotdm*	IoT Data-centric Middleware	tcpmd5	TCP-MD5 Support library
l2switch	Separate Layer 2 switching	toolkit	Quickstart Toolkit
lacp*	Link Aggregation Control Protocol	tpf*	Topology Processing Framework
lispflowmapping	LISP Mapping Service	tsdr*	Time Series Data Repository
nic*	Network Intent Composition	ttp	TTP Project
openflowjava	OpenFlow Protocol Library	usc*	Unified Secure Channel
openflowplugin	OpenFlow Plugin	vtn	VTN (Virtual Tenant Network)
		yangtools	YANG Tools

\* → New in Lithium release



# 4<sup>th</sup> Release "Beryllium" Production-Ready Open SDN Platform

Graphical User Interface Application and Toolkit (DLUX / NeXT UI)

AAA AuthN Filter

OpenDaylight APIs REST/RESTCONF/NETCONF/AMQP

### Base Network

- Host Tracker
- L2 Switch
- OpenFlow Forwarding Rules Mgr
- OpenFlow Stats Manager
- OpenFlow Switch Manager
- Topology Processing

### Enhanced Network Services

- |                                |                            |                             |
|--------------------------------|----------------------------|-----------------------------|
| AAA                            | Messaging 4Transport       | SNMP4SDN                    |
| Centinel – Streaming Data Hdlr | NetIDE                     | Time Series Data Repository |
| Controller Shield              | Neutron Northbound         | Unified Secure Channel Mgr  |
| Dev Discovery, ID & Drvr Mgmt  | OVSDB Neutron              | User Network Interface Mgr  |
| DOCSIS Abstraction             | SDN Integration Aggregator | Virtual Private Network     |
| Link Aggregation Ctl Protocol  | Service Function Chaining  | Virtual Tenant Network Mgr. |
| LISP Service                   |                            |                             |

### Network Abstractions (Policy/Intent)

- ALTO Protocol Manager
- Fabric as a Service
- Group Based Policy Service
- NEMO
- Network Intent Composition

Controller Platform  
Services/Applications

Data Store (Config & Operational)

Service Abstraction Layer/Core

Messaging (Notifications / RPCs)

- |                      |           |        |           |       |      |       |         |         |      |       |      |       |                |       |            |
|----------------------|-----------|--------|-----------|-------|------|-------|---------|---------|------|-------|------|-------|----------------|-------|------------|
| OpenFlow 1.0 1.3 TTP | OF-Config | OVSD B | NETCONF F | LIS P | BG P | PCE P | CAPWA P | OPFLE X | SX P | SN MP | US C | SN BI | IoT Http/CoA P | LAC P | PCMM /COPS |
|----------------------|-----------|--------|-----------|-------|------|-------|---------|---------|------|-------|------|-------|----------------|-------|------------|

Southbound Interfaces & Protocol Plugins

OpenFlow Enabled Devices



Open vSwitches



Additional Virtual & Physical Devices



Data Plane Elements (Virtual Switches, Physical Device Interfaces)



# What does it mean to program in ODL?

1. Get familiar with the Model-View-Control approach for app development in a modular fashion
  - YANG (Yet Another Next Generation) **Model** for data, RPC and notifications
  - RESTconf **View** generated automatically
  - **Implementation** in Java to handle data changes, notifications and RPC call backs
2. Get familiar with platform essentials (maven, config subsystem, dependencies) and useful tools
3. Learn about existing projects and reuse modules
  - No need to change code of other projects. Just link them as binary libraries

# YANG modeling

# YANG

▶ Data modeling language that is also the preferred configuration language for NETCONF protocol

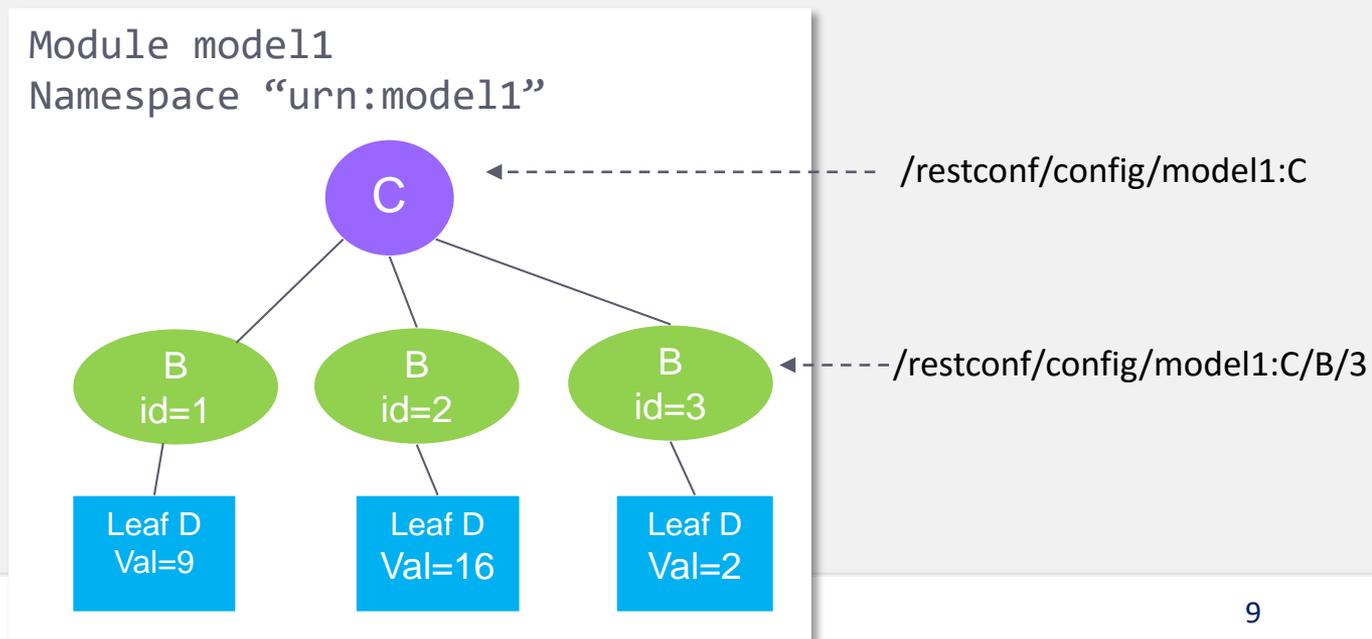
▶ Further reads:

- [YANG introductory tutorial](#)
- [RFC 6020 - YANG - A data modeling language for NETCONF](#)

```
module model 1 {  
  
    namespace "urn:model 1";  
    prefix model 1;  
    yang-version 1;  
  
    revision 2015-04-06 {  
        description "Initial revision"  
    }  
  
    grouping A {  
        list B {  
            key id;  
            leaf id {  
                type uint32;  
            }  
            leaf D {  
                type uint32;  
            }  
        }  
    }  
  
    container C {  
        uses A;  
    }  
}
```

# MD-SAL Data Access

- ▶ Model-driven SAL is the kernel of the OpenDaylight controller!
- ▶ It manages the contracts and state exchanges between every application. It does this adaptation by managing centralized state
- ▶ Takes in the YANG model at runtime and constructs the tree in the data store
  - For the YANG model in previous slide, here is the view of the root and its children



# MD-SAL Data Access (contd.)

- ▶ When the model is compiled with maven, you will see classes generated in Model1Data.java, B.java, and C.java
- ▶ InstanceIdentifier is used as a pointer to a child. Following points to the first child node in the figure

```
InstanceIdentifier iid = InstanceIdentifier.builder(C.class)
    .child(B.class, new BKey((long)1))
    .build();
```

- ▶ ReadOnlyTransaction, and WriteTransaction are used to access the data store.

```
B updatedB = new BBuilder().setD((long)19).build();
WriteTransaction modification = dataBroker.newWriteOnlyTransaction();
modification.merge(LogicalDataStoreType.CONFIGURATION, iid, updatedB, true);
modification.submit();
```

- ▶ Transaction can also be batched

# Two types of data store

- ▶ For the same model, there are two types of data store maintained by MD-SAL
  - **Config store:** App developers typically use this to store user input and associated derived state for apps
  - **Operational store:** Typically used to keep transient ephemeral state
- ▶ Choice of store has implications on RESTconf and persistence
  - Config store is always kept persistent. There is control on how many replicas to keep
  - Operational store cannot be changed over RESTconf

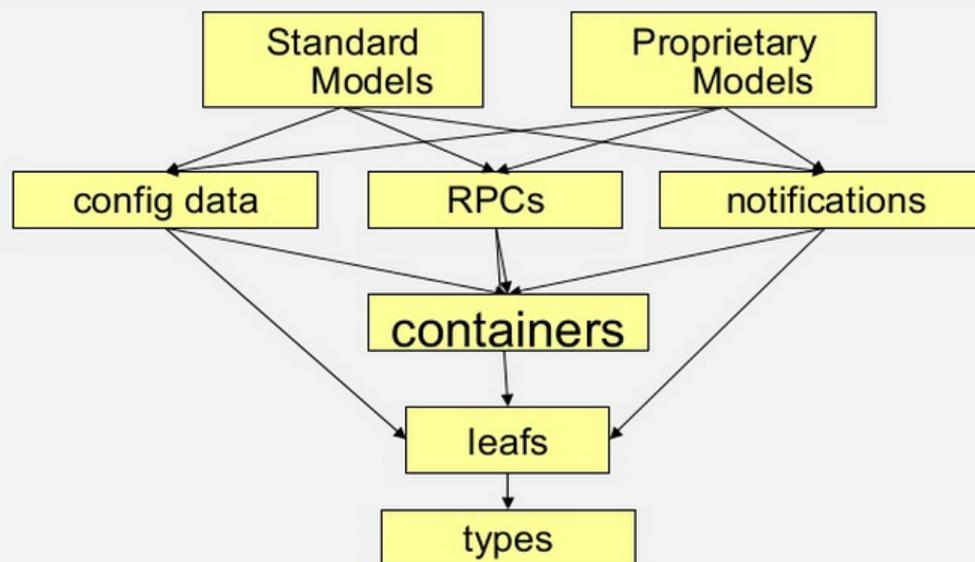
# YANG not restricted to Just Data Store

## ► Notifications:

- Publish one or more notifications to registered listeners

## ► RPC:

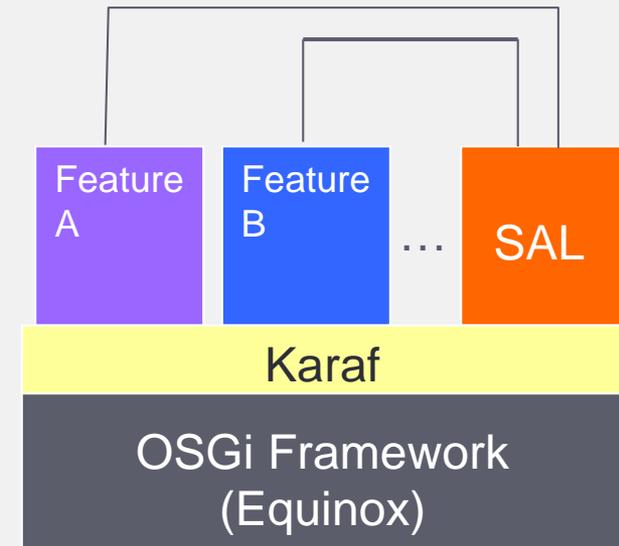
- Perform procedure call with input/output, without worrying about actual provider for that procedure



# Poking into the basic platform

# Java, Interface, Maven, OSGi, Karaf

- ▶ Java chosen as an enterprise-grade, cross-platform compatible language
- ▶ Java Interfaces are used for event listening, specifications and forming patterns
- ▶ Maven – build system for Java
- ▶ OSGi:
  - Allows dynamically loading bundles
  - Allows registering dependencies and services exported
  - For exchanging information across bundles
- ▶ Karaf: Light-weight Runtime for loading modules/bundles
  - OSGi based. Primary distribution mechanism for Helium



# REST APIs

- ▶ OpenDaylight has significant support for REST APIs
- ▶ Restconf allows for checking config and operational state
  - `feature:install odl-restconf`
  - [http://localhost:8181/restconf/....](http://localhost:8181/restconf/...)
- ▶ List of exposed Northbound APIs are autogenerated using swagger.
  - `feature:install odl-mdsal-apidocs`
  - <http://localhost:8181/apidoc/explorer>



## OpenDaylight RestConf API Documentation

Controller Resources

Mounted Resources

Below are the list of APIs supported by the Controller.

config(2013-04-05)	Show/Hide
config-logging(2013-07-16)	Show/Hide
flow-capable-transaction(2013-11-03)	Show/Hide
flow-topology-discovery(2013-08-19)	Show/Hide
ietf-netconf-monitoring(2010-10-04)	Show/Hide
kitchen-service-impl(2014-01-31)	Show/Hide
network-topology(2013-07-12)	Show/Hide
network-topology(2013-10-21)	Show/Hide
opendaylight-action-types(2013-11-12)	Show/Hide
opendaylight-flow-statistics(2013-08-19)	Show/Hide
opendaylight-flow-table-statistics(2013-12-15)	Show/Hide
opendaylight-group-statistics(2013-11-11)	Show/Hide
opendaylight-inventory(2013-08-19)	Show/Hide
POST /config/	
GET /config/opendaylight-inventory:nodes/	
PUT /config/opendaylight-inventory:nodes/	
DELETE /config/opendaylight-inventory:nodes/	
POST /config/opendaylight-inventory:nodes/	
GET /config/opendaylight-inventory:nodes/node/{id}/	

# ODL's opendaylight-inventory.yang (Lithium)

```
module opendaylight-inventory {
  namespace
  "urn:opendaylight:inventory";
  prefix inv;

  typedef node-id {
    type inet:uri;
  }

  typedef node-connector-id {
    type inet:uri;
  }

  typedef node-ref {
    type instance-identifier;
  }

  typedef node-connector-ref {
    type instance-identifier;
  }
}
```

```
grouping node {
  leaf id {
    type node-id;
  }

  list "node-connector" {
    key "id";
    uses node-connector;
  }
}

grouping node-connector {
  leaf id {
    type node-connector-id;
  }
}

container nodes {
  list node {
    key "id";
    uses node;
    . . .
  }
}
```

Augmentation made by  
OpenFlow plugin for storing:

1. All switch description
2. All OpenFlow features
3. All tables and its flows
4. All meters, groups

```
notification node-updated {
  status deprecated;

  leaf node-ref {
    type node-ref;
  }
  uses node;
}

notification node-connector-updated {
  status deprecated;

  leaf node-connector-ref {
    type node-connector-ref;
  }
  uses node-connector;
}

notification node-removed {
  status deprecated;

  leaf node-ref {
    type node-ref;
  }
}

notification node-connector-removed {
  status deprecated;

  leaf node-connector-ref {
    type node-connector-ref;
  }
}
```

# ODL's Inventory Config Data Store

<http://localhost:8181/restconf/config/opendaylight-inventory:nodes>

```
▼<nodes xmlns="urn:opendaylight:inventory">
  ▼<node>
    <id>controller-config</id>
  </node>
  ▼<node>
    <id>openflow:1</id>
    ▼<table xmlns="urn:opendaylight:flow:inventory">
      <id>0</id>
      ▼<flow>
        <id>561183150</id>
        ▼<match>
          ▼<ethernet-match>
            ▼<ethernet-destination>
              <address>00:00:00:00:00:02</address>
            </ethernet-destination>
            ▼<ethernet-source>
              <address>00:00:00:00:00:01</address>
            </ethernet-source>
          </ethernet-match>
        </match>
        <flags/>
        <flow-name>mac2mac</flow-name>
        <priority>512</priority>
        ▼<instructions>
          ▼<instruction>
            <order>0</order>
            ▼<apply-actions>
              ▼<action>
                <order>0</order>
                ▼<output-action>
                  <output-node-connector>openflow:1:1</output-node-conr
                  <max-length>65535</max-length>
                </output-action>
              </action>
            </apply-actions>
          </instruction>
        </instructions>
      </flow>
    </table>
  </node>
</nodes>
```

```
▼<flow>
  <id>0</id>
  <match/>
  <flags/>
  <flow-name>allPacketsToCtrl</flow-name>
  <priority>0</priority>
  ▼<instructions>
    ▼<instruction>
      <order>0</order>
      ▼<apply-actions>
        ▼<action>
          <order>0</order>
          ▼<output-action>
            <output-node-connector>CONTROLLER</output
            <max-length>65535</max-length>
          </output-action>
        </action>
      </apply-actions>
    </instruction>
  </instructions>
  <idle-timeout>0</idle-timeout>
  <hard-timeout>0</hard-timeout>
  <table_id>0</table_id>
  <buffer_id>0</buffer_id>
</flow>
```

```
▼<flow>
  <id>1
  ▼<matc
    ▼<et
  </e
  </mat
  <flag
  <flow
  <prio
  ▼<inst
    ▼<in
  </i
  </ins
  <idle
  <hard
  <tabl
  <cook
  <buft
  </flow
  </table
```

# ODL's network-topology.yang (Beryllium)

<http://localhost:8181/restconf/operational/network-topology:network-topology>

```
module network-topology {
  namespace
  "urn:TBD:params:xml:ns:yang:network-topology";

  typedef topology-ref {
    type leafref {
      path "/network-
topology/topology/inet:uri";
    }
  }

  typedef node-ref {
    type leafref {
      path "/network-
topology/topology/node/inet:uri";
    }
  }

  typedef link-ref {
    type leafref {
      path "/network-
topology/topology/link/inet:uri";
    }
  }

  grouping tp-attributes {
    leaf tp-id {
      type inet:uri;
    }
    leaf-list tp-ref {
      type tp-ref;
    }
  }
}
```

```
  grouping node-attributes {
    leaf node-id {
      type inet:uri;
    }
    list supporting-node {
      key "node-ref";
      leaf node-ref {
        type node-ref;
      }
    }
  }

  grouping link-attributes {
    leaf link-id {
      type inet:uri;
    }
    container source {
      leaf source-node {
        mandatory true;
        type node-ref;
      }
      leaf source-tp {
        type tp-ref;
      }
    }
    container destination {
      leaf dest-node {
        mandatory true;
        type node-ref;
      }
      leaf dest-tp {
        type tp-ref;
      }
    }
    list supporting-link {
      key "link-ref";
      leaf link-ref {
        type link-ref;
      }
    }
  }
}
```

```
  container network-topology {
    list topology {
      key "topology-id";
      leaf topology-id {
        type inet:uri;
      }
      container topology-types {
      }
      list underlay-topology {
        key "topology-ref";
        leaf topology-ref {
          type topology-ref;
        }
      }

      list node {
        key "node-id";
        uses node-attributes;
        list termination-point {
          key "tp-id";
          uses tp-attributes;
        }
      }

      list link {
        key "link-id";
        uses link-attributes;
      }
    }
  }
}
```

# Extension to models through Augments

- ▶ Unless specified otherwise, YANG models can be augmented in a new namespace to add extra data within an existing data model.
- ▶ For example, odl-l2-switch augments above inventory to store hosts learned

```
import opendaylight-inventory { prefix inv; revision-date 2013-08-19; }

augment "/inv:nodes/inv:node/inv:node-connector" {
    ext: augment-identifier "address-capable-node-connector";
    uses address-node-connector;
}
```

```
curl http://localhost:8080/restconf/operational/opendaylight-
inventory:nodes/node/openflow:1/node-connector/openflow:1:1
{ "node-connector":
  [ { "id": "openflow:1:1",
    "address-tracker:addresses": [ {
      "address-tracker:last-seen": 1404918398057,
      "address-tracker:ip": "10.0.0.1",
      "address-tracker:first-seen": 1404918392926,
      "address-tracker:mac": "f2:0c:dd:80:9f:f7"
    } ]
  } ]
}
```

# Code Design

- ▶ yangtools generates data model classes
  - Builders for generating immutable objects also generated
  - identity mapped to interfaces
  - All else are mapped to classes
  - Identity is the only one that allows inheritance
  
- ▶ Developers have to write the gateway adapters
  - Converting from one data model to another
  - Translating between a non-ODL source and ODL app

[https://wiki.opendaylight.org/view/YANG\\_Tools:YANG\\_to\\_Java\\_Mapping](https://wiki.opendaylight.org/view/YANG_Tools:YANG_to_Java_Mapping)

# Data Store Management

## ▶ Persistence

- Everything that gets into a shard is stored in-memory and on the disk.
- Restarting the controller will reconstitute the state of the shards from the persisted data.

## ▶ Clustering

- Multiple controller instances can interwork and share state (for backup reasons rather than load-sharing)
- One leader (writer) and multiple followers (read-only) where state consistency is achieved using RAFT algorithm

## ▶ Sharding

- The big in-memory MD-SAL tree is broken up into a bunch of smaller sub-trees such that 1 model is 1 shard (e.g., inventory, topology and default).

# Southbound Plugins: OpenFlow and NETCONF

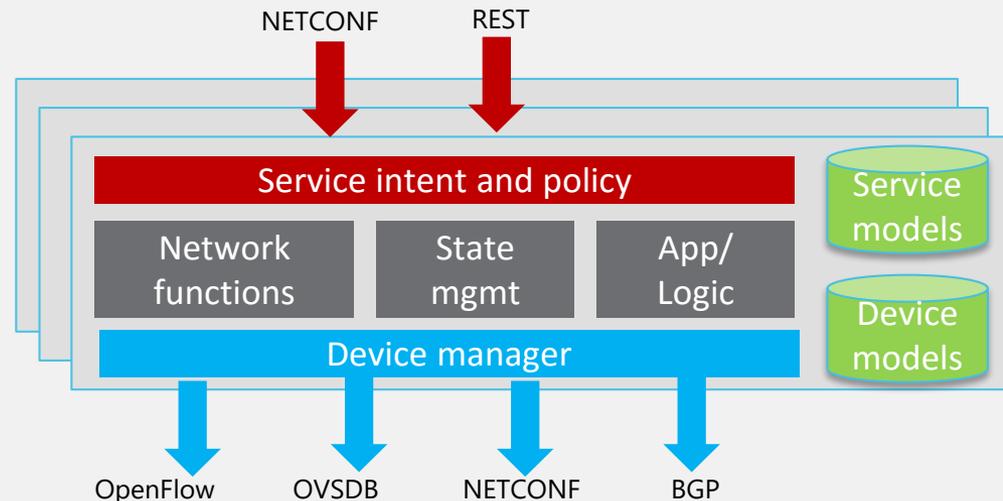
# OpenDaylight External Interfaces

## ▶ Southbound:

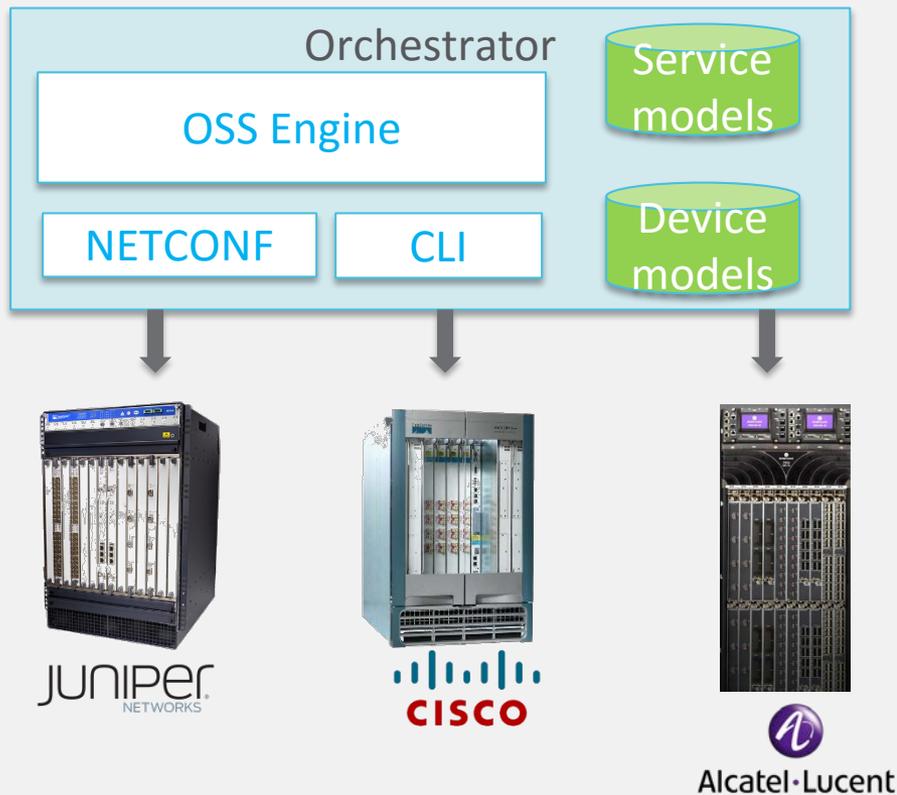
- Built-in NETCONF client and BGP speaker to interact with legacy control planes and orchestrate them
- Plugin for OpenFlow 1.0 and 1.3 support
- Plugin for OVSDB that works with OVS schema (VTEP schema in future)

## ▶ Northbound:

- REST and NETCONF for receiving intent data



# NETCONF based Multi-Vendor OSS



Models are vendor specific and still need vendor specific adapters in the orchestrator layer. This makes case for OpenConfig

# Questions?