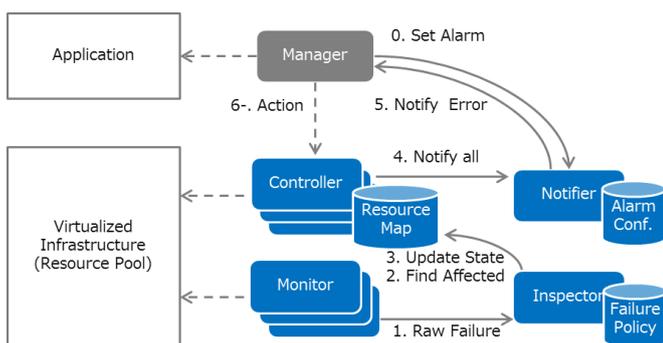


# Doctor: Failure Detection and Notification for NFV

▶ The Doctor project in OPNFV is implementing immediate detection and notification of resource failures to application managers.

▶ Telecom nodes, due to their stringent high availability requirements, often come in an Active-Standby (ACT-STBY) redundant configuration. When virtualized, the manager of such a virtualized node cannot directly access the state of the underlying infrastructure; thus, it requires fault notification on the ACT node application to be able to instantly switch to the STBY node application. This reduces the downtime to nearly zero for such telecom nodes hosting thousands of mobile connections.

▶ Our solution monitors a resource pool by an external resource monitor. As soon as one of the Monitors detects a failure in any of the resource pool elements, the Inspector will analyze the events based on its failure policies/templates and will update the status of the resource. The Notifier will also send an immediate notification to the Application Manager of the affected node, which in turn switches to the STBY node application.



▶ In addition to last year's OPNFV Summit the demo was extended to show, e.g.

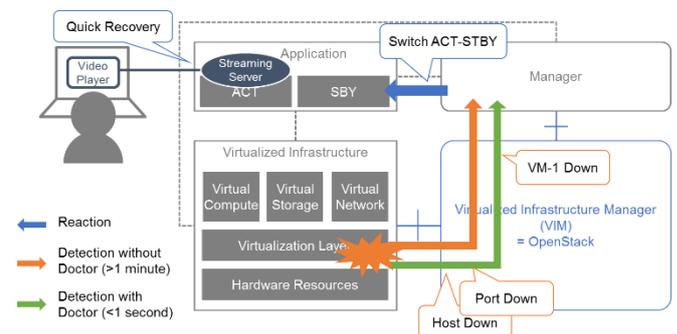
- Integration of Congress / Vitrage as Inspector for policy based failure identification & analyzes.
- Integration of additional Monitors such as Nagios and Collectd to extend the range of monitored events and demonstrate Doctor's pluggable architecture.
- Extension to include both Nova and Neutron.

▶ Doctor team members have completed blueprints for Aodh, Nova, and Congress (as of June 2016).

Project	Blueprint	Spec Drafter	Developer
Aodh	Event Alarm Evaluator	Ryota Mibu <sup>1</sup>	Ryota Mibu <sup>1</sup>
	New nova API call to mark nova-compute down	Tomi Juvonen <sup>2</sup>	Roman Dobosz <sup>3</sup>
Nova	Support forcing service down	Tomi Juvonen <sup>2</sup>	Carlos Goncalves <sup>1</sup>
	Get valid server state*	Tomi Juvonen <sup>2</sup>	Tomi Juvonen <sup>2</sup>
	Add notification for service state change*	Balazs Gibizer <sup>4</sup>	Balazs Gibizer <sup>4</sup>
Congress	Push Type DataSource Driver*	Masahito Muroi <sup>5</sup>	Masahito Muroi <sup>5</sup>
	Adds Doctor Driver*	Masahito Muroi <sup>5</sup>	Masahito Muroi <sup>5</sup>

\*NEW since last PoC (Nov'15)      <sup>1</sup>NEC <sup>2</sup>Nokia <sup>3</sup>Intel <sup>4</sup>Ericsson <sup>5</sup>NTT

▶ In OpenStack without our solution, the delay to notify the application manager is in the order of several minutes. Given such long notification delay, in a telecom scenario, e.g. thousands of mobile subscribers will be disconnected from their cellular network. In contrast to that, in our solution, we can perform such failure notification within 1 second. In the demo, we show how well Doctor performs in failure detection and notification, thus being able to meet the high availability requirements of telecom node applications in NFV-based virtualized network systems.



@ OPNFV Summit 2015, OpenStack Summit Austin



@ OPNFV Summit 2016

▶ Visit Doctor Wiki: <https://wiki.opnfv.org/doctor>

# Doctor: Failure Detection and Notification for NFV



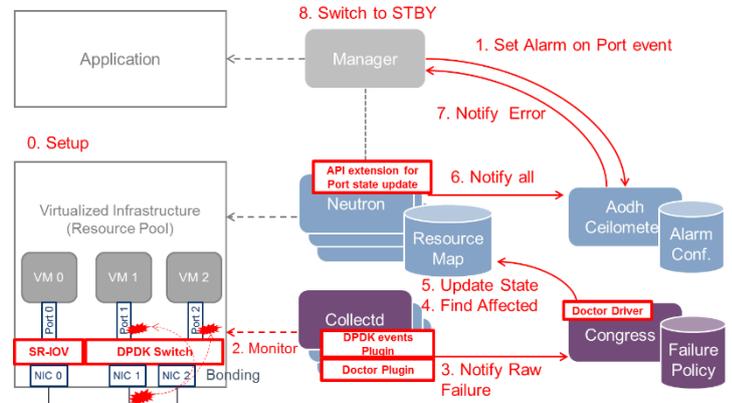
## OPNFV Doctor + SFQM + OpenStack Congress



➤ **Congress** is the OpenStack Policy Evaluation Engine for enforcing **flexible and dynamic failure identification policies** based on the demands of the Operator.

**DPDK** is a set of open-source libraries and drivers for fast packet processing.

**Collectd** is a daemon which periodically gathers performance statistics about the system it is running on. Those statistics can then be used to find current performance bottlenecks and predict future system load.



- Doctor Inspector is realized by Congress:
  - Congress collects various information in the cloud.
  - Congress can calculate a mapping from raw failure events to logical failure events using this information.
  - Congress can take an action following policy rules defined by the Administrator.

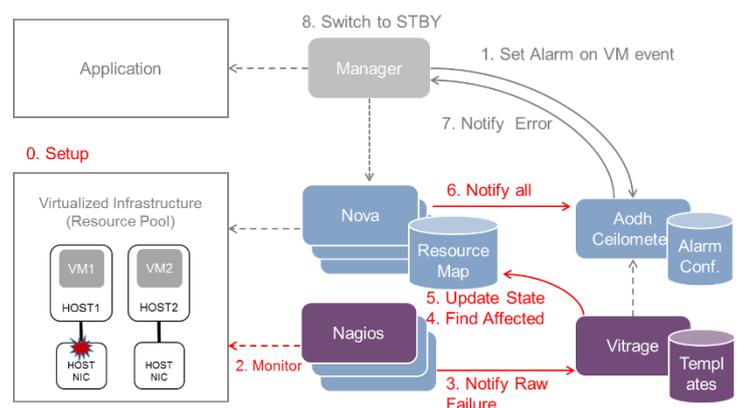
- The demo shows:
  - Event-driven notification with DPDK stats and collectd.
  - Failure mapping in logical resource view.
  - API extension for Neutron.
  - End-to-end toolchain from physical failure, all Doctor building blocks, up to the Application Manager level.

## OPNFV Doctor + OpenStack Vitrage



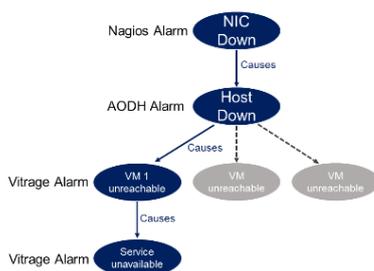
➤ **Vitrage** is the OpenStack Root Cause Analysis Project for analyzing and expanding alarms and events, yielding insights regarding the **root cause of the problems**, and **deducing the existence of faults** before they are directly detected.

- Doctor Inspector is realized by Vitrage:
  - Vitrage has push and pull interfaces to various monitoring tools (e.g. Nagios, Zabbix) and to OpenStack.
  - Vitrage can be configured using templates to execute actions e.g. raise deduced alarm/state.
  - Vitrage can provide root cause indicators to the Application Manager.



- Demo steps as follows:
  - Vitrage receives alarm from Nagios once host NIC fails.
  - Vitrage finds affected resources based on its cloud topology graph and templates.
  - Vitrage calls Nova to mark host down and update the states of the affected VM(s).
  - Nova notifies Aodh which raises an alarm, Application Manager gets alarm notification and transfers control to VM2 (STBY).

Vitrage Demo - Root Cause Analysis



Vitrage States View

