

HPA Enhancements (For Casablanca Release)

Model-driven HPA

The "model-driven" in the title above refers to a data model for definition of VNF descriptors, as specified by the ETSI NFV SOL001 specification. Starting with Casablanca, the term "onboarding model" may also be used in reference to this data model. The data model is based on a TOSCA Simple YAML 1.1 /1.2 language encoding. In the future (post-Casablanca) this data model may also be encoded using YANG. The data model is persisted as an artifact in the TOSCA CSAR file (zip archive) as specified by the ETSI NFV SOL004 specification. The specification of HPA requirements is incorporated in the VNFD definition.

~~In R2, HPA requirements for each VDU (VNFC) of VNF are represented as a policy. Though policy-driven HPA will continue to be there in R3, our intention is to auto-create HPA policies from VDU model of the service. Intention is to avoid creation of policies outside of model. Since TOSCA, industry standard, is chosen to represent model, even VNF vendors can provide the VNF HPA requirements.~~

~~Model driven HPA needs following enhancements:~~

- ~~• Creation of HPA policies dynamically from the VNFD/VDUD~~
- ~~• Verification of HPA requirements defined in the VNFD/VDUD during service creation~~

Model-driven HPA User Story

- VNF developers create a VNFD as part of the VNF package. This may be done by hand or using visual design tools. Once the package has been created, it is validated using the VNFS SDK tools, or by other equivalent means.
- VNF developers use the VNFS SDK tools will validate the structure and integrity of the CSAR file and the syntactic/semantic validity of the VNFD, including the HPA requirements.
- ONAP operators onboard VNFs into ONAP. As part of the onboarding process, the VNF undergoes a set of operator specific acceptance tests. If the VNF acceptance is successful, the VNF becomes part of the design time catalog.
- As part of VNF instantiation, the VNFD contents is propagated to the Policy Framework.
- The Policy Framework translates the HPA requirements specified within the VNFD into OOF placement policies.
- The OOF subsequently uses these policies to determine proper homing and placement of VNF components.

Model-driven HPA Work Items

- Validation of HPA requirements contained in the VNFD
- On-boarding of VNFDs into SDC
- Propagation of VNFDs from SDC to Policy
- Translation of VNFD contained HPA requirements into OOF policies

Note - It is assumed that no changes to OOF and SO are required from the model-driven HPA perspective.

Affected Projects

- Modeling Subcommittee
 - Endorse the R3 resource data model (HPA included)
- VNFS SDK
 - Verify the HPA model semantically and syntactically
- SDC :
 - Ensuring that HPA additions do not affect the functionality
 - Propagation of VNFDs to POLICY
- POLICY framework
 - ~~Look for new models~~
 - ~~Download models~~
 - ~~Get hold of HPA information~~
 - Translation of HPA requirements into OOF policies

Prerequisites:

- ~~SO implementing TOSCA models~~
- Adoption of TOSCA VNFDs across all use cases
- Creation of TOSCA models for existing use cases (vFW/vDNS and vCPE)

SRIOV NIC Support

HPA feature in OOF facilitates placing VDUs on the site that has compute nodes supporting needed SRIOV-NIC cards. Operators like to have flexibility of placing the workloads on non SRIOV-NIC machines in case there are no matching profiles. In HEAT templates, there is a challenge to have this kind of flexibility. HEAT templates are normally created for SRIOV-NIC as design time as there is special HEAT logic for SRIOV-NICs. If non-SRIOV site is chosen at run time, then this special HEAT information needs to be replaced dynamically with normal vSwitch information based on the results from OOF.

- SO/Multi-Cloud

- Check the profile returned by OOF.
- For each CP of VDUs/VNF,
 - Check whether that CP requires SRIOV-NIC and verify whether the NIC supported by that profile.
- If above condition is true for all CPs, then there is no change required in HEAT information.
- If not, modify appropriate portion of HEAT information from SRIOV-NIC to normal vswitch.

VF-C with HPA policies

In ONAP, there are two projects are doing similar orchestration - SO and VF-C. SO to OOF integration with HPA is done in R2. VF-C to OOF integration with HPA is slated for R3.

We believe there are following activities:

- Though VF-C understand TOSCA based models, we will not touch existing HPA capability support as VF-C will be moving to support new TOSCA models (R3 based models). Hence, we don't expect to work on auto-creation of HPA policies based on R2 DM HPA information. Only R3 DM HPA to HPA policies will be supported. This work will not be started until VF-C supports R3 model.
- Support VF-C and OOF integration with manually created HPA policies
 - VF-C (GVNFM) to OOF integration to get the best region and set of flavors (compute profiles)
 - Pass this compute profile (for each PDU) to the Multi-Cloud instead of asking Multi-Cloud to create the compute profile.
- Since, GVNFM is not used by any use case, make vFW use case work with VF-C.

Affinity/Anti-Affinity & Impact on HPA compute profile selection

There is some discussion about introducing Affinity and Anti-Affinity policies in OOF as part of R3. As we understand, affinity/anti-affinity rule granularity is at

- NFVI-PoP level.
- Availability zone level
- NFVI-node level

NFVI-node level may not be relevant for ONAP.

Currently, input to HPA filter is just set of regions (NFVI-PoPs). This may need to be enhanced to support "availability zones".

Also, HPA filter may need to output best availability zone within NFVI-PoP (region).

Features that are being carried over from R2

OOF Enhancements

R2 OOF returns the VNF placement information by giving best cloud region and set of compute profiles in that cloud-region for constituent VDUs of the VNF. R2 OOF does this by matching the operator defined policies with the capabilities of cloud-regions. Some features could not be implemented in R2. Intention is to implement following in R3

- Logging & statistics and Visualization
 - Ability for deployment administrator to view the compute profiles used by VNFs over time (historical data)
 - VNFs that could not be placed due to non-matching compute profiles.
 - VNFs that are placed, but the optimal compute profile could not be found.
 - VNFs that are placed with various scores.
- Best site selection with respect to compute profiles: R2 has ability to select the best compute flavor on per site basis. That is, HPA filter on per input candidate site, checks whether there is any matching compute profile. If there is one, it stops searching on rest of candidates. In R3, intention is to search through all input candidate sites, get the best compute profile in each site and select the site whose compute profile (combined) score is highest.
- Ensuring that results are consistent irrespective order of the filter execution: Currently, HPA filter is run last. If HPA filter is run first, the result should be same. There are few enhancements required in OOF execution order of filters. It may be that, there are two phases of execution. First phase requires selection of candidate list and second phase is to select the best of all the selected sites.

SO Enhancements

HPA functionality is tested with vCPE use case in R2. As we understand, SO has various workflows that could be different for various use cases. In R3, intention is to ensure that

- HPA compute profile selection for vFW and vDNS use cases.
- ??

Multi-Cloud Enhancements

- Adding more HPA features (e.g OVS DPDK)
- ???