VOLTHA-GO-CONTROLLER

Radisys Corporation - CONFIDENTIAL

ONOS vs VGC Deployment

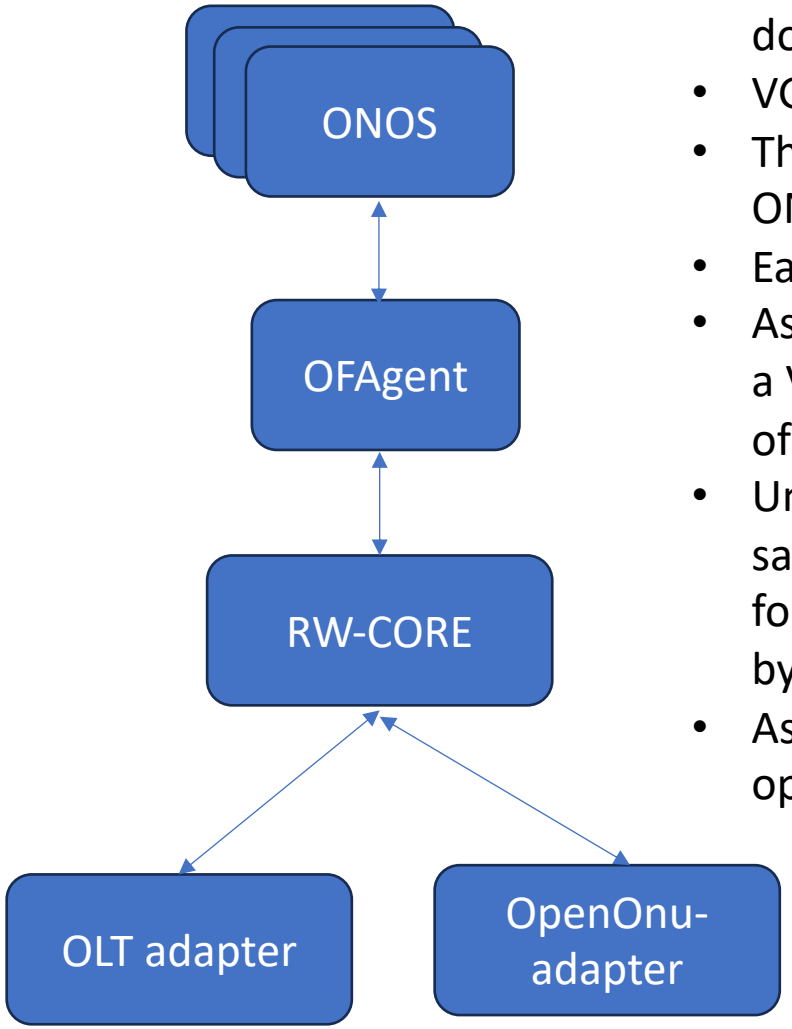

- VGC does the same functionality as is being done by ONOS and OFagent combined.
- VGC directly interfaces with the rw-core.
- The SDN apps in VGC are static whereas the ONOS apps can be dynamically loaded
- Each VOLTHA stack has once instance of VGC
- As once instance of VGC is deployed in a VOLTHA stack, it scales along with the number of stacks.
- Unlike ONOS, VGC does not run as a cluster. The same philosophy of fail/recover fast, as is used for other VOLTHA components is also employed by VGC.
- As VGC supports in-transaction resiliency , operating in a clustered mode is not a necessity.

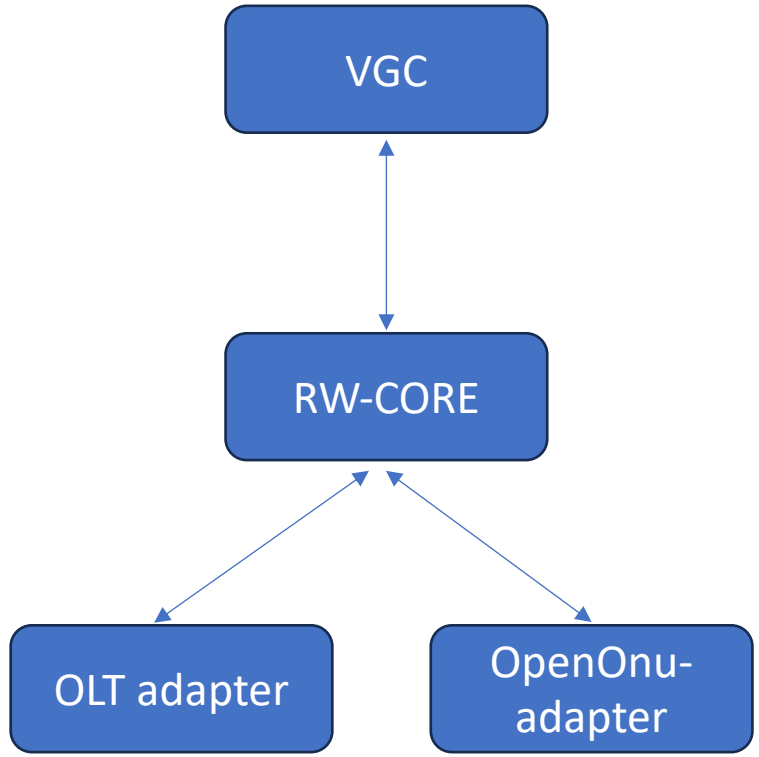

ONOS Vs VGC Resource consumption

Radisys

ONOS

Setup	Berlin 151 Setup		
Config	vStacks used: 4 OLTS per vStack : 2 PON ports per OLT : 16 ONUs per PON : 32	vStacks used: 6 OLTS per vStack : 2 PON ports per OLT : 16 ONUs per PON : 32	vStacks used: 8 OLTS per vStack : 2 PON ports per OLT : 16 ONUs per PON : 32
Num ONUs	4096	6144	8192
			CPU (cores)
total	18.24	21.61	24.38
onos+atomix	8.18	9.59	11.22
etcd	3.59	4.21	4.03
kafka	0.49	0.58	0.51
vStacks	4.6	6.44	7.62
			Memory (GB)
total	22404.93	26870.67	29826.35
onos+atomix	14467.59	15805.76	16103.59
etcd	1367.93	1773.33	2045.56
kafka	1674.23	1610.27	1666.97
vStacks	4797.61	7049.1	9331.02
			Time taken (secs)
ONUs PORT-UP	272	415	517

VGC

Setup	Berlin 141 setup		
Config	vStacks: 2 OLTs per vStack: 8 ONUs per OLT: 512	vStacks: 3 OLTs per vStack: 11 ONUs per OLT: 512	vStacks: 3 OLTs per vStack: 16 ONUs per OLT: 512
#ONUs	8192	16896	24576
		CPU(Cores)	
etcd	8.35	9.26	9.66
kafka	0.53	0.32	0.59
vStacks	7.28	12.04	20.48
TOTAL	16.16	21.62	30.73
		MEMORY(MB)	
etcd	2225.01	3210	4140.44
kafka	2904.44	2912.58	3024.44
vStacks	11487.56	21899.95	31506.15
TOTAL	16617.01	28022.53	38671.03
		Time-taken(seconds)	
ONUs PORT-UP	144	295	484

ONOS Vs VGC Resource consumption

Radisys

- **ONOS & Atomix (with 12k Subs)**

Memory (MB)	CPU
18453	12.82

- **VGC with 24576 Subs (3 vStacks, 16OLTs/vStack, 512 ONUs/OLT)**

vstack Instance	Memory (MB)	CPU
vstack1-voltha-go-controller	243.21	0.660
vstack2-voltha-go-controller	203.02	0.650
vstack3-voltha-go-controller	203.89	0.620
Total	**650.12**	**1.930**

- **VGC uses 1/10th of the CPU used by ONOS and atomix**
- **VGC uses 1/3rd of the memory used by ONOS and atomix**

(One more point to note is the CPU and memory comparison is done with 25ksubs for VGC and only 12k subs with ONOS)


Radisys
Thank You

Resource consumption

	CPU (millicpu/millicores)			MEMORY (MB)		
POD_name	**AVG**	**MIN**	**MAX**	**AVG**	**MIN**	**MAX**
vstack1-voltha-go-controller	183.23438	91	479	29.890625	23	62
vstack2-voltha-go-controller	191.29688	93	360	33.59375	26	65
vstack3-voltha-go-controller	180.17188	74	440	29.59375	22	68
vstack4-voltha-go-controller	176.625	90	462	31.390625	25	70
vstack5-voltha-go-controller	172.76563	88	361	30.3125	22	74
TOTAL	904.09377mi	436mi	2102mi	154.78125MB	118MB	339 MB

Resource Consumption for ~10K subs, 5 Vstacks, 4 OLTs per Stack, 512 ONUs per OLT

- **Max CPU consumption per VGC instance is 479 milicpu**
- **Max memory consumed by an instance is 74 MB.**

- **All the tests done here are with persistency enabled.**
- **Considerable amount of CPU and memory saved by replacing ONOS with VGC.**