OpenVMS Deployment Comparison & VSI Cloud

Abstract

In this article, we will explore the 5 most common ways to deploy VSI OpenVMS x86-64 - both in the Cloud and on-prem - and how each approach performs in a multitude of dimensions generally considered by IT decision makers and system architects: billing, CapEx vs OpEx, availability, reaction time, outsourcing responsibility, scalability, etc.

Why Consider VSI Cloud

For decades, OpenVMS has been the gold standard for reliability, security, and uptime, and that has been achieved all on the back of bare metal deployments. For 40+ years running on physical hardware was not just the best way to achieve that level of stability: it was the only way.

As technology evolves and new concepts and technologies hit the market, so does the opportunity to merge different paradigms to create a solution with the best of multiple worlds. And the Cloud is one of these brave new worlds, with its own promises but also challenges. For many, the idea of moving OpenVMS to the Cloud feels like trading a proven fortress for a promise of flexibility, but what if the Cloud isn't a discussion of a *step-down/step-up*, but rather a complementary way forward?

The prototyping nature of early Cloud generations are many years in the past, with nowadays maturity and reliability taken as an undisputed given. This shift is owed mostly to modern Cloud platforms' distributed architectures, redundancy, and automated failover.

Beyond reliability, the Cloud brings other advantages that more and more VSI customers are looking for:

1. Reduced costs: much lower CapEx, lower OpEx
2. Simplified billing
3. High availability technologies
4. Geographical unpinning
5. Performance tuning on the fly
6. Full official 24/7 support
7. Easy deployment, easy reproducibility

In this article, we will explore what our research has shown to be the 5 most common ways to deploy VSI OpenVMS - both in the Cloud and on-prem - and how each performs.

The 5 ways to have your OpenVMS

Our market analysis and customer interactions show that there are 5 main ways to deploy the OpenVMS operating system on X86:

1. **Hypervisor on physical server on Prem**
    a. You buy a server, install a supported hypervisor, and run it in your own datacenter. This hardware is owned and maintained by you.
2. **Hypervisor on physical server on Outsourced datacenter**
    a. Similar to the above, but the server is now running on a datacenter rented space, i.e. not owned by your company. This hardware is owned by you, but it may be maintained by you or by the datacenter.
3. **Hypervisor on physical server on Cloud**
    a. You are "renting" a physical server, which resides in a datacenter, both owned and physically maintained by your Cloud provider. Compared to physical server on premises, the main difference here is that you are not responsible for any of the hardware in your system, only the software.
4. **Generic Cloud - Virtualized**
    a. You have access to a virtual machine, which utilizes a slice of a physical server or group of servers.
    b. You are not pinned to a given specific server: the exact hardware utilized may change at any moment's notice.
    c. Note: it may be either *a nested virtualized Cloud* solution or a *non-nested (native) virtualized Cloud* solution.
5. **VSI Cloud - Virtualized**
    a. Similar to above, but you are utilizing a Cloud service provided by VSI that has been certified and tuned to work seamlessly with the VSI OpenVMS operating system

The comparison dimensions

CapEx vs OpEx

One of the first bi-dimensional dispersions we may delve into when analyzing the 5 different ways of deploying the OpenVMS operating system is CapEx vs OpEx.

It is a fact that the exact expenditure and the ratio between these 2 will vary greatly depending on the characteristics of your system, so exact quantitative USD quotes may not be a great way to characterize the 5 deployments, but rather a qualitative approach is needed:

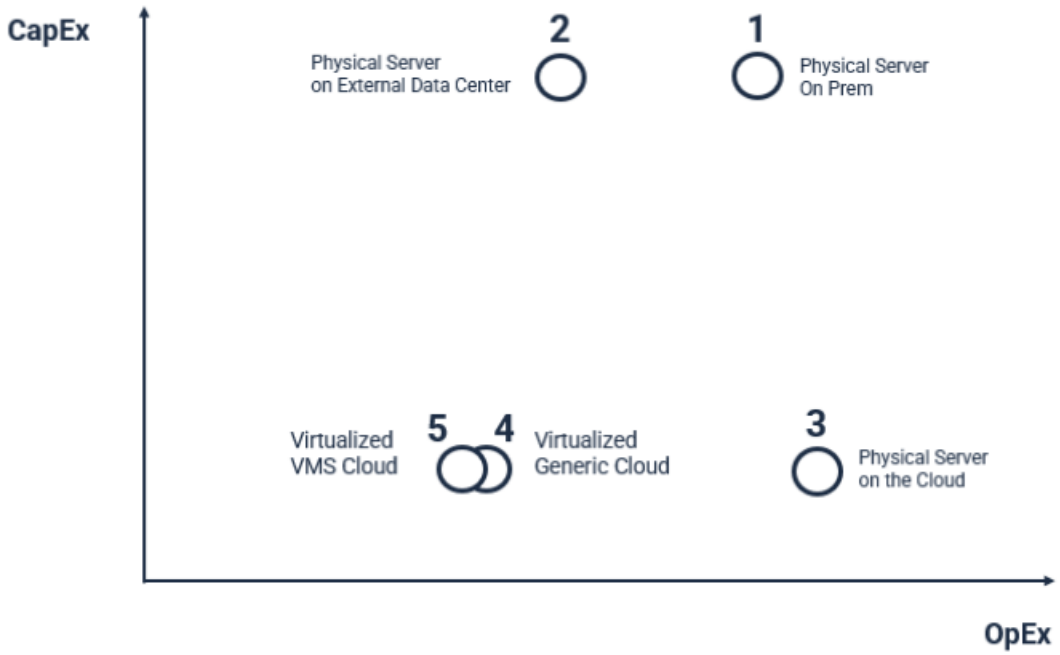

The reason for this dispersion is that deployment methods 1 and 2 require the upfront acquisition of a server, thus elevated CapEx, with method 1 in particular also requiring further ancillary equipment acquisitions, such as power and cooling infrastructure. Deployment method 3 does not require such upfront acquisition, but its cost is instead spread over the usage of the server rental period, thus increasing OpEx significantly higher than any other option.

For deployment methods 4 and 5, on the other hand, one does not have any CapEx other than account start-up costs and initial system setup (engineering time), which are orders of magnitude smaller than server acquisition cost. Additionally, by not renting an exclusive server as in method 3, but only a small slice of it, the OpEx is still small, which is seen as a major advantage by providing financial flexibility and simplifying enterprise financial planning.

Reaction time vs Outsourcing responsibilities

The balance between the “reaction time for hardware issues” vs “outsourcing responsibilities” is an important consideration when comparing the 5 deployment models.

On-premise setups allow for immediate, direct intervention when problems occur, provided the organization has the necessary in-house expertise and resources available at all times. The reaction time is therefore as fast as your own organization can provide. In contrast, outsourced server hosting shifts much of the technical responsibility to external providers, which can improve overall reliability but may lead to slower response times due to reliance on external support teams and agreed service levels. Cloud bare metal hosting goes a step further by automating much of the infrastructure and offering scalable support, but it can introduce even more delays when addressing complex issues, as users must depend on the provider’s support framework and limited system-level access. In essence, when dealing with physical servers, the more operational control is outsourced, the less direct oversight an organization retains, and consequently also the reaction time to hardware issues may be slower, due to more stringent SLAs being generally cost

prohibitive or simply not available. There is therefore a clear trade-off between faster hands-on response vs reduced day-to-day management burden.

On the other hand, when dealing with virtualized servers in the Cloud, provided your Cloud solution has innate software-based live migration of a running virtual machine, hardware issues are no longer as common, as servers are more often rotated for preventive maintenance, load sharing and replacement of degrading or failing components. In this case, while an original server is being prepared to go on maintenance, the infrastructure is already shifting its computational load to other performant hardware units, without any need of user intervention. Although the fix or improvement to the specific hardware server under maintenance may take many hours, from the perspective of the computational workload the reaction time is in fact nearing instantaneous, as most live migration technologies do so in less than a few milliseconds. Therefore, in the context of Cloud providers with live migration - even though you are outsourcing most of your platform and hardware responsibilities - you still get an extremely fast reaction time to hardware maintenance downtimes, even when running isolated non clustered single instances. The conclusions above are better understood via the qualitative scatter below:

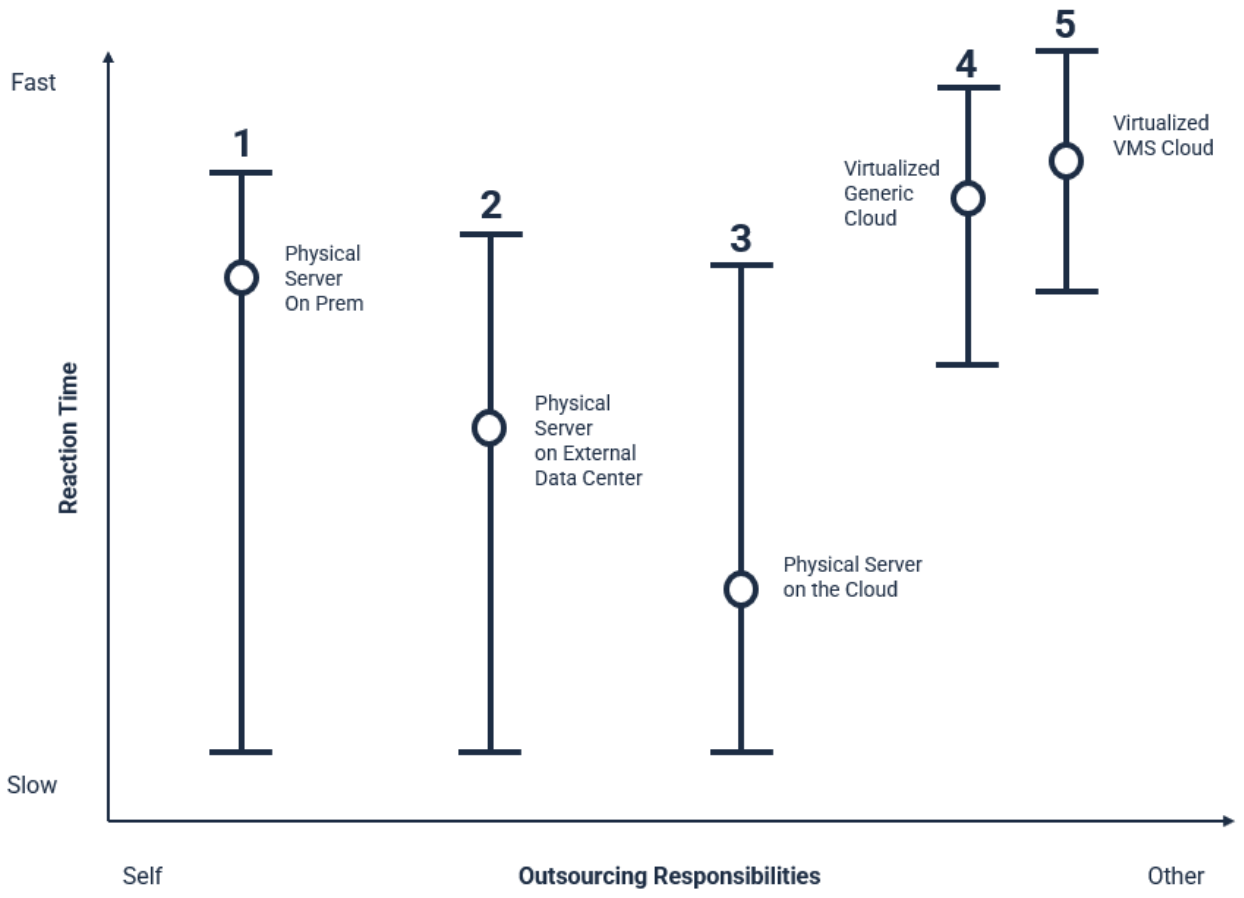

Economies of scale and gains from standardization

Cloud services benefit from economies of scale that typically outweigh those of on-premises deployments, offering cost and operational advantages for most organizations. In this domain, size does matter: large scale enables more efficient use of resources, bulk purchasing, quicker adoption of new technologies, etc. In contrast, on-premises setups often involve high upfront investment, dedicated IT teams, and underutilized systems maintained for occasional peak demand, not to mention a reduced negotiation power when dealing with utilities providers, equipment resellers, and real estate costs.

VSI Cloud offers all the economies of scale advantages from the big Cloud providers by directly leveraging their infrastructure, but at the same time adding OpenVMS-based standardization to this deployment. It's having the cake and eating it too: your OpenVMS deployment is backed by state-of-the-art datacenters with millions of square feet of area, but you are still able to spin up a new OpenVMS instance at will, with ease, and by fully documented and automated means supported by VSI.

The usage of such standardized means guarantees a much smoother future transition to ancillary services, such as VSI Managed Services. Should the need arise, it is now possible within short notice of one of your key employee's retirement to hand over full responsibility of your production deployment to VSI, with the knowledge that the VSI personnel taking over is fully trained in compliance with VSI Cloud console access, MFA and the multitude of security tools and protocols in place.

Further comparison dimensions

As can be seen in the previous 3 chapters, the 5 deployment methods can be analyzed from different perspectives, and since each customer gives differing weights to each perspective as part of their strategic deployment decisions, there is no single-fits-all winner between the 5 options.

To get a better understanding of these perspectives, the table below is provided with shaded colors accounting for the qualitative advantages of each of the 5 deployment methods over a range of perspectives that may be of essence to a customer decision on how to deploy their OpenVMS system.

<table>
<tr><th></th><th>Category</th><th>Physical server On Prem</th><th>Physical server on External Datacenter</th><th>Physical server in The Cloud</th><th>Generic Cloud</th><th>VSI Cloud</th></tr>
<tr><td>A</td><td>CapEx</td><td colspan="2">Initial Server costs</td><td colspan="3">No Initial server costs</td></tr>
<tr><td>B</td><td>OpEx</td><td>Utilities, spares, personnel</td><td>Datacenter costs, spares</td><td>Monthly server costs</td><td colspan="2">Virtualized servers are generally 1 order of magnitude cheaper than a comparable exclusive server in the Cloud</td></tr>
<tr><td>C</td><td>Reliability/</td><td colspan="3"></td><td colspan="2"></td></tr>
</table>

	Availability	Unless v-motion OR live migration is enabled, and multiple servers are running concurrently, the breakdown of a single machine means a downed node			Physical server downtime is not felt by the instance as a live migration is done within milliseconds	
D	**Fast Reaction Time**	Limited by your team 24/7 availability and distance to your own datacenter	Limited by your datacenter provider capabilities and SLA	Highly limited for small to medium customers due to worse SLAs	If live migration is intrinsic, only massive datacenter outages may affect your deployment. Single hardware failures are nearly instantaneously swapped out by a replacing host	
E	**Billing**	Many separate bills, employee costs	Many separate bills	Separate bills from your Cloud provider vs your license and support from VSI		Unified single bill from VSI
F	**Future Proof**	We guarantee compatibility of newer OpenVMS releases with a limited hardware list certified by the hypervisor		We cannot guarantee new OpenVMS releases compatibility as the range of hardware configurations is too broad		We guarantee compatibility with newer OpenVMS releases
G	**Official Support**	Support engineers may not be familiar with the peculiarities of your environment		Support engineers may not be acquainted with other Cloud providers and their hypervisor software configuration		VSI Support highly experienced in this deployment
H	**Geo Mobility**	Physical moving of server with potentially days to weeks downtime if no transitional nodes are added		Hours to days downtime, convoluted process	Very small downtime, generally in the range of minutes, easy to do and undo	
I	**Gains from standardization / Ancillary services**	Hardest access to managed services given non standardized connection methods	Generally more standardized ways of access are enforced by datacenter provider	Easier access to managed services		
J	**Managed Services**	As long as certified hardware is used, it will be easy for VSI to provide services as we know how the platforms work seamlessly		A different Cloud provider may mean VSI needs more time to get acquainted to your custom setup, and out tools may not be compatible		Extensive knowledge, experience and set of standardized tools
K	**Vertical Scalability**	Server upgrades require downtime	Server upgrades downtime and longer waiting times to commence	Server upgrades may not even be possible	Server upgrades with downtime limited to a reboot.	
L	**Elasticity**	Performance is static as per original dimensioning			Performance can be adjusted up and down as needed, generally in real-time	
M	**Ease of Deployment**	Requires physical server setup, long waiting times		Reduced waiting times, but requires configuration of the server underlaying settings	Quick to configure a new deployment, but no official documentation on the platform given by VSI	Quick to configure, official VSI documentation present, articles with examples

N	**Reproducibility**	High likelihood of manual mistakes or differences from one node or machine to the next. Harder to audit each step		Requires custom scripting made by the customer		Fully unattended automated scripts, removing people in the loop, prevent manual mistakes when creating a new node
O	**Environment**	Less efficient cooling and energy usage, wastage of hardware, unused bandwidth and performance	Commercial datacenters have more efficient cooling and electricity usage due to economies of scale	Your hardware server is dimensioned to the peak demands, thus wastage	Your computing runs on a much more optimized platform, in a server at a higher utilization and thus lower wastage	
P	**Latency**	Reduce distance to your end point application	Medium distance to your end point application	Increased distance to your endpoint application		
Q	**Full control of data ownership**	You have direct physical access and overview of your data	You have overview of your data but maybe at times no direct physical access	No direct physical access to your data, but still full overview		
R	**Usability in physical process control**	Reduced latency, capability of operating with most protocols and connection methods	Increased latency, potential protocol and communication issues	Latency and limited support for custom protocols and custom interfaces		
S	**Supplier lock-in**	No lock-in: change in server vendor, network and storage devices possible	Small lock-in: moving of datacenters will cause downtime if not lean executed	Cloud provider lock-in, generally to single infrastructure provider		Lock into cloud service, not necessarily to infrastructure provider

-

Conclusion

At the end of the day, there's no universal answer to whether OpenVMS should stay on a hypervisor on a physical server or move to the Cloud: each option has its trade-offs. Some organizations will prioritize the control of on-prem systems, others will lean into the flexibility and scalability the Cloud offers – but many will find themselves somewhere in between, in a true mash-up style of on-prem, off-prem and Cloud virtualized hybrid clusters. In true logical Boolean style, VSI Cloud is thus not a matter of "OR", but rather of an "AND", and an AND in full symbiosis with the time proven OpenVMS technology for distributed hybrid resilient clusters.

What really matters is finding the right fit for your team, your goals, and the peculiarities of your application. For the cases where the right fit may be difficult to decide or to precisely tailor, please do count on VSI's expertise to assist you with the decision-making process.

Whatever path you choose, one thing is certain: VMS Software is here to help you protect your investment in OpenVMS, now and into the future.