SUNLIGHT

faster ■ better ■ smarter

Use Case

SUNLIGHT

faster ■ better ■ smarter

Use Case

Creating the most cost-effective way for enterprises to host their high-performance workloads

How MSPs and Datacenters can deliver a differentiated alternative to deploying in public cloud or on-premise for enterprise customers

Overview

Enterprises are constantly under pressure to reduce the costs of IT, yet at the same time their IT needs are only increasing. Traditional enterprise workloads, such as databases and ERP systems are put under more load as business fully digitise. New, highly strategic workloads such as big data, artificial intelligence and machine learning create whole new demands which have to be accomodated.

The virtualisation technology that underpins the majority of these workloads introduces significant performance overheads. This either necessitates major over-investment in physical infrastructure, or having to maintain specific workloads on bare-metal, with associated reduced manageability and resilience.

Simply adding more servers and storage to a datacenter in the hope of chasing Moore's Law is no longer seen as an option. Instead, enterprises have to increase efficiency and flexibility. In many cases, public cloud infrastructure such as Amazon Web Services, Microsoft Azure and Google Compute Engine are seen as the obvious way to go. However the hidden costs of guaranteeing performance in the public cloud often leads to an uncontrollable TCO - leading to disruptive reversals of cloud migrations.

What if there was a way for MSPs and datacenter operators to offer their enterprise customers a cost-effective way to host demanding workloads with guaranteed high performance at a fraction of the cost of the mainstream public cloud operators?

Running a high performance workload in the datacenter

As high-performance workloads become more prevalent, the need for infrastructure that can support very high IO and low latency increases. Fortunately technologies like 100Gb Ethernet, NVMe flash storage and NUMA memory have emerged to meet the need. However, traditional virtualization stacks - VMware and HyperV - are simply not architected to be able to support the levels of performance that these technologies offer. The virtualisation tax these hypervisors introduce is up to 75-90% of IOPS performance in the case of NVMe storage . This necessitates the build out of significantly more infrastructure than would be required if the workload ran on bare-metal.

Leaving the workload running on bare metal is one solution, however most enterprises are minimising the number of applications running in this mode wherever possible. Running on bare metal not only increases management costs, but reduces resilience and availability - as workloads and associated services can't easily fail over, and it reduces flexibility - as workloads become tied to specific hardware configurations.

Running a high performance workload in the public cloud

Perhaps the answer to this conundrum is to run these workloads in the public cloud? Indeed many enterprises are choosing to go 'cloud first' - where the starting assumption is that a workload will be deployed in the cloud, and the case for deploying on-premise must be argued strenuously.

For lower performance workloads - this strategy can work just fine. However for high performance workloads, it is often a recipe for huge unbudgeted costs and even failure. In fact IHS Markit found that 74% of IT directors questioned had had to reverse a cloud migration - with half of those reversals due to performance or security issues .

Public cloud infrastructure typically comprises hosted services running on shared, multi-tenant machines. There are inherent security risks associated with applications running on shared CPU and memory systems as witnessed by the slew of CPU flaws (Spectre, Meltdown etc..) that are regularly discovered and announced. Ensuring isolation in public cloud requires investment in dedicated bare metal servers and recreating an equivalent virtualized multi-tenant infrastructure. This requires personnel investment or expensive licensing for platforms such as VMware cloud in AWS, Azure or GCE.

Running applications fully virtualized in the public cloud introduces severe performance penalties, similar to those described above for running on-premise. These include IOPS limits on instance types - e.g. an 80k IOPS limit for AWS instances, and inconsistent and low performance storage necessitating the use of expensive guaranteed performance options like AWS Provisioned IOPS (PIOPS). This often far eclipses the anticipated savings from moving to the cloud in the first place - cloud computing becomes expensive quickly, especially when trying to implement high performance workloads.

In addition, instances are generally offered as fixed units of consumption - specific memory, CPU, networking and storage sizes that don't match the workload's actual requirements. This makes it necessary to overprovision, and pay for, certain resources. For example to achieve 80K IOPS on an instance in AWS, the maximum number of cores on the physical host have to be subscribed - even if they are not required. This results in wasted infrastructure costs, and sometimes excessive licencing costs for workloads licenced by memory or CPU cores.

1 Sunlight's own research - see https://performance.sunlight.io

2 IHS Markit "The Bi-Directional Cloud Highway" - 2019

Does hyperconverged infrastructure provide the answer?

Hyperconverged infrastructure delivers similar scale-out capability to public cloud but with better TCO. While in the cloud, new infrastructure can be added on-demand, for HCI, infrastructure can be added a box at a time.

Whereas guaranteed performance in the public cloud is expensive and often an unbudgeted cost, hyperconverged infrastructure makes it easy to scale as infrastructure can be added building block at a time - each HCI node containing compute, storage and networking. These resources are automatically added to the overall pool, and can be instantly put to use. For predictable workloads that don't experience huge peaks and troughs of usage - this model can work extremely efficiently.

Unfortunately, to-date most hyperconverged stacks are still based on VMware or KVM - and experience the same performance penalties discussed previously, so whereas scaling out is simple, it is still extremely inefficient.

Introducing NexVisor - the hyperconverged stack with bare metal performance.

Sunlight offers a complete hyperconverged stack, including the NexVisor hypervisor, software defined storage, software defined networking and a management dashboard. Sunlight has been rearchitected from the ground-up to remove the virtualisation overhead of traditional hypervisors. Sunlight allows workloads to run at near-bare metal performance - unlocking the full investment in hardware. For demanding workloads that require high IO and low latency, Sunlight is able to perform an order of magnitude faster than any other hypervisor on the market.

In the following example, Sunlight is able to run an Oracle workload at over 3x the executes-per-second rate of VMware or the AWS native hypervisor, and with a tiny fraction of the read latency.

Sunlight allows virtualized workloads to be deployed at the lowest possible TCO.

Comparing Sunlight, VMware and AWS Native performance for Oracle

ORACLE® DATABASE

Oracle DB performance metric - Executes (SQL) per second

Proportional load with increasing number of cores.

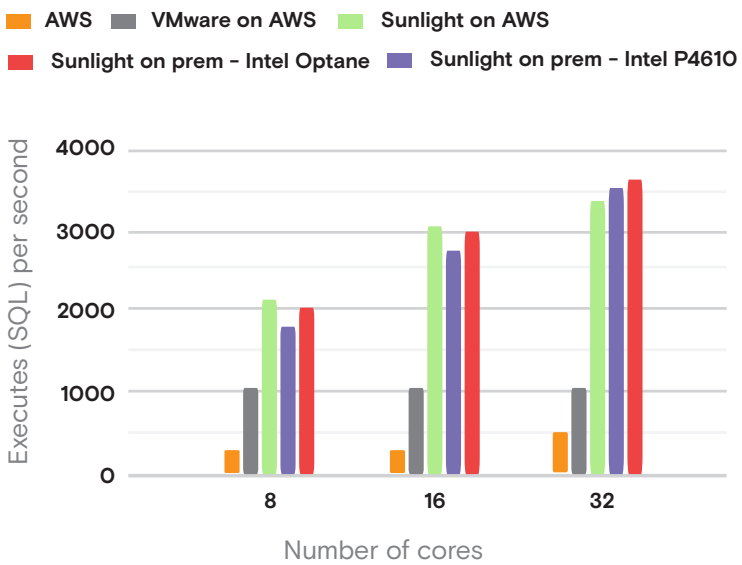

Oracle DB performance metric - Read latency

Proportional load with increasing number of cores.

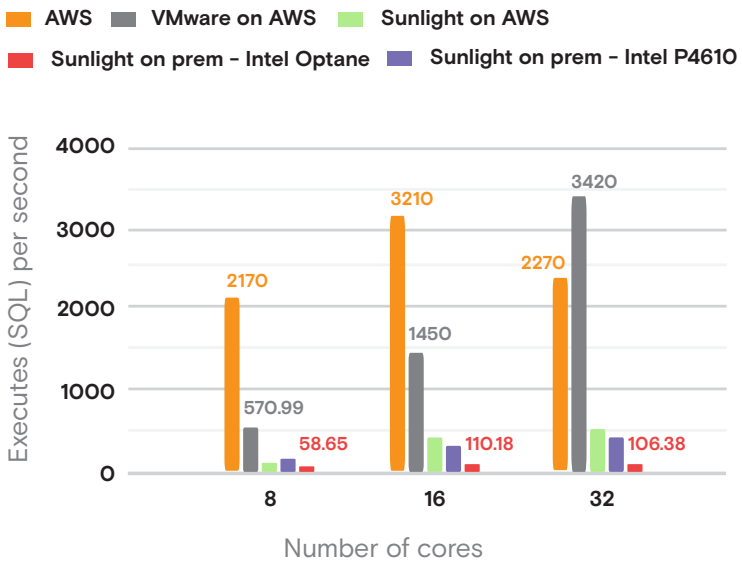

The best of all worlds - hosted Sunlight private cloud

Offering enterprise customers a hosted Sunlight private cloud cluster delivers a very attractive alternative to the customer building out the infrastructure in their own datacenter. The key benefits are:

- True bare metal IO performance - 10x higher than running on traditional virtualisation technology with over 1 Million storage IOPs per VM instance as standard.
- Dedicated server clusters provide complete security guarantees and transparency about exactly which workloads are running on which CPU cores and accessing which physical memory
- Virtualisation manageability & high availability, including enterprise class features such as High availability, complete VM instance lifecycle support
- Cloud flexibility and opex pricing
- Option to deploy in a hybrid on-premise plus hosted configuration

For the service provider, Sunlight hyperconvergence simplifies the cost and dependencies on external technologies such as expensive SAN disk arrays, and custom network infrastructure such as firewalls, fibre channel and infiniband. In addition, Sunlight's flexible monthly pricing options make it easier to offer opex pricing options to the customer.

The impact of Sunlight on datacenter costs

Monthly Costs

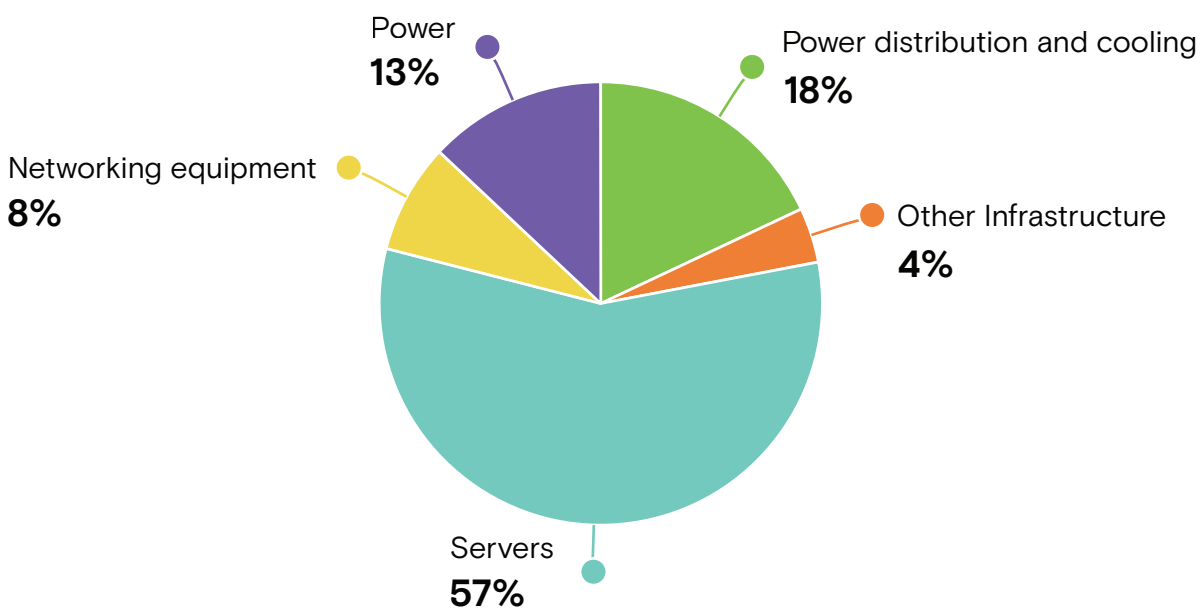

3 year server and 10 year infrastructure amortization

When datacenter costs are pulled apart and amortized , we find that the major cost is the servers themselves, followed by power and cooling for those servers. It is clear that minimising the server estate is the best way to reduce overall costs.

Costs of deploying a high-performance private cloud: VMware vs Sunlight

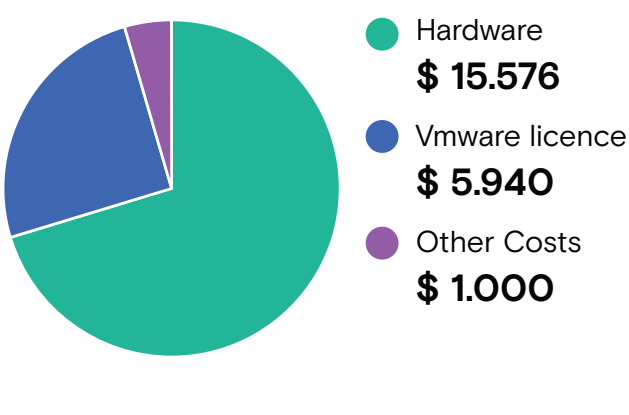

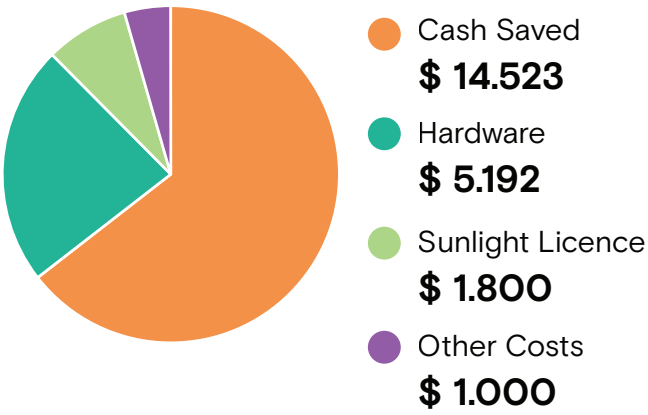

vmware®

SUNLIGHT

In the above example, for a high-performance workload, a six-node VMware cluster can be replaced with a 2-node Sunlight cluster, still offering similar or better performance for IO-bound workloads, at a 64% infrastructure cost saving. This enables the service provider to offer an extremely competitive alternative to the customer deploying on-premise or in the public cloud.

[3] James Hamilton - https://perspectives.mvdirona.com/2010/09/overall-data-center-costs/

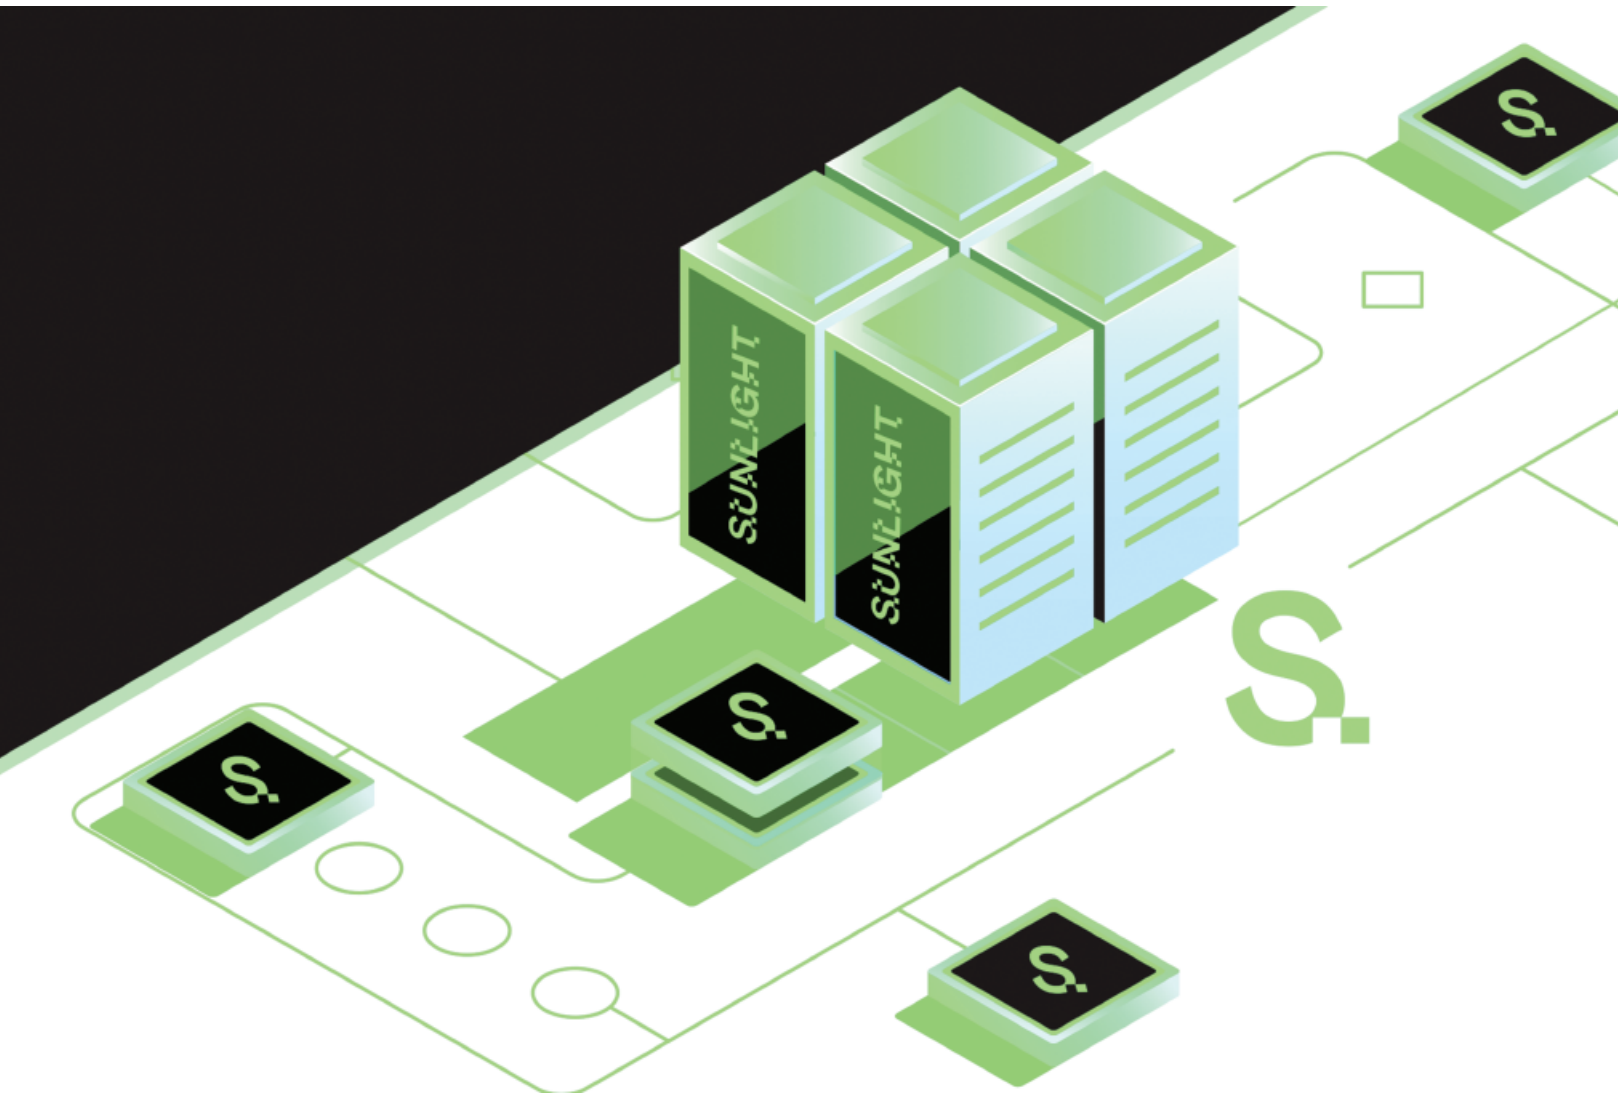

Conclusion

With a private cloud hosted solution built on Sunlight technology, Datacenter operators can provide the most cost effective and high performance solution to their enterprise customers. This enables them to compete with the Hyperscaler service offerings with superior performance and managed costs.

Get in touch today to learn more about how Sunlight solutions can help you grow your datacenter business.

sales@sunlight.io +1-888-477-0009 Castle Park, Cambridge, UK