



Some like it hot – VM and container migration in hybrid cloud environments

Nine considerations for maximizing flexibility and utilization

Table of Contents

Getting started – what should IT departments consider when architecting for VM and container migration? 2

Making sense of workload placement 2

Intel® Optane™ DC persistent memory increases VM density, TCO and uptime 4

Intel provides the architecture for the hybrid cloud. 4

Processing virtualized workloads with Intel® Xeon® Scalable processors 4

Expanding memory capacity beyond DRAM limits with Intel® Optane™ SSDs 5

Improving time-to-deployment with Intel® Select Solutions 5

Conclusion 5

Picture the scene. A smooth-running IT organization seamlessly migrates its virtual machine (VM) workloads across the globe, as utilization rises or falls in its various sites. As the staff go home in Oregon, and the server loads fall, then deep learning inference workloads spin up to take advantage of the spare capacity. As the staff in the London office trickle in to start the day, the financial modeling applications that have been running overnight power down, and move to Tokyo to continue running as the office there closes for the evening. Perhaps some of the workloads will be transitioning to and from public cloud environments as the economics and capacity constraints allow.

Or another scene. A smaller company, struggling with constrained resources and ever-increasing workload demands, wraps vital legacy apps in containers to postpone their technical debt, probably to have those apps eventually re-architected as cloud-native microservices when the development resources are available.

Or again, a company running high-performance computing (HPC) workloads exports the workload to the public cloud for testing and tuning. This allows the company to test and tune, without the need to build and support the test hardware environment. After optimizing the workload, it can be brought back on-premise to be applied to the latency-critical massive data sets that are used in their production environment.

All these scenarios require smooth and efficient container and VM migration between and among public and private clouds. How does all this happen? What do IT organizations need to consider to make this a reality? What are the architectural and hardware decisions that are important to achieve this flexibility and efficiency?

This paper is about the considerations IT organizations need to address to make container migration and consolidation efficient and practical, the hardware choices that underpin it, and the technical resources – hardware, software, industry enablement, and advice – which Intel provides to make it possible.

Getting started – what should IT departments consider when architecting for VM and container migration?

Here are the nine key considerations which need to be worked through to support the various scenarios above.

1. Hot or cold migration?

Do you intend to shutdown workloads and spin them up somewhere else as a permanent move, effectively replicating the workload on a different cloud? Or do you, as in the first scenario above, intend to 'hot' migrate a workload? Cold migration of some workloads can have economic and practical benefits. Effectively, VMs are moved and data replicated, and the initial instances shut down. Applications where this can be useful include server consolidation, legacy application preservation or simplifying the management of many workloads into one virtualized support model. Cold migration can also be used to evacuate older legacy system platforms into a newer virtualized model whether it is on-premise in private clouds, or off-premise in the public cloud. This helps to reduce the impact of data center legacy debt by removing older, obsolete hardware and moving workloads into modern infrastructure.

'Hot migration' is where the VM is moved live between sites, either with its data or with dual access to a low-latency network or data architecture. The VM and its workloads continue to run. This is useful in the 'follow the sun' model outlined in the first scenario above, or where workloads are being moved between public and private clouds as capacity and priorities require. The scenario includes such applications as deep learning, inference, modeling and other compute-intensive apps. A key benefit to hot migration is having a well-developed application architecture that allows for your data to be highly accessible with minimal downtime. The VMs or containers that are migrated are generally used as 'workers' in the environment. These worker VMs use only the data required for the immediate tasks at hand, but the primary data continues to be centralized for flexible functionality. Hot migration does require modernized cloud functions like orchestration and modeling to effectively migrate workloads.

Hot and cold migrations have fundamentally different requirements on management, network, storage and processing environments. Working out which workloads will be moved where, why, and how frequently is the first and primary step of a sound migration strategy.

Making sense of workload placement

Evaluating workload placement and creating a holistic multi-cloud strategy is complex and nuanced. However, the rewards are worthwhile. Considering the business needs and the technical characteristics of each workload—performance, security, integration, and data requirements – is the first step. Then conducting a full total cost of ownership (TCO) analysis will deliver a better understanding of how different combinations (modernize on premises, integrate public cloud for strategic workloads, etc.) will deliver the most cost-effective mix of solutions for your business. Lastly, evaluating the cloud ecosystem – suitable cloud service providers (CSPs) aligned to business needs, software-as-a-service (SaaS) maturity and commercially available solutions for particular workloads, and the amount of personnel with cloud expertise that an organization can or cannot tap into.

Intel has conducted research with hundreds of major enterprises to see how they undertake workload placement across hybrid- and multi-cloud environments. That research, abstracted alongside the experience of Intel's system architects and illustrated with real customer examples, is available as the paper '[Optimal Workload Placement in a Multi-Cloud World: Public, Hybrid, and Private Clouds](#)'.

The paper also includes a step-by-step guide to working through the TCO and other key considerations of where to place workloads and why.

2. Workload placement

Working through a rigorous workload placement methodology can help to clarify decisions and optimize utilization and total cost of ownership (TCO). Do you have an economic and practical model to establish which workloads should be placed where on your hybrid multi-cloud architecture, and those which it makes sense to migrate? You can determine optimal workload placement by considering things like latency, security, data size and integration requirements of the applications and workloads you need to run. See the box-out for more details and links.

3. Management infrastructure

Are you moving towards HCI? How are you going to oversee the management and automate the provision of virtualized services? Are you using a 'single pane of glass' management infrastructure so that you can see the performance of all your servers and systems?

The more comprehensive the management infrastructure, the easier it is to move VMs, maximize utilization, and propagate changes. A complex migration strategy will require clarity of oversight and decision-making. Enterprises need to ensure that they have the management tools and orchestration capability to see the utilization levels of the different components of their server estate. For many, there are multiple systems running various workloads more as an accident of history than as the output of a systematic plan to maximize the utilization of their estate. The first stage in achieving this is having a 'single pane of glass' view of the entire IT infrastructure. Management requires decision, and decisions can be coded into profiles or policies. Decisions made on a regular basis can be automated and scripted, helping your management infrastructure to potentially manage itself. This allows administrators to work on improvements, while menial tasks are self-administered through software.

4. Data architecture

For many workloads, data characteristics become a significant issue when the application is migrated. For applications, accessing massive data sets then porting or replicating the data can be impractical. And, when using public cloud, further elements of economic complexity need to be factored in. Yet many applications need to run close to their data. For some, it may be practical to port or to synchronize the data along with the VM, while for others the data may need to sit in a low-latency repository accessible by all the sites running the VM. Public cloud service providers offer you the ability to create profiles for different types of workload requirements and manages data access for you. When architecting the private or hybrid cloud solutions, the same items must be addressed. Intel® architecture implements telemetry which feeds into modern hybrid cloud models to help prioritize utilization and balance workloads in the cloud.

Additionally, the General Data Protection Regulation (GDPR) and other data governance regulations need to be understood and applied as data moves between sites and countries. With GDPR it is also essential that data collections reside in governed regions, and demonstrable efforts can be proven to ensure the privacy of customer data.

5. Security

It is vital to keep data and VMs secure within and between environments. Maximizing network and data security, and data persistence, is critical in managing successful migrations as the VMs and data flow north-south (in and out of individual systems across the network). Protecting data as it flows in and out of the system is greatly enhanced by using software-defined networking (SDN) technology. The latest SDN technologies combine very robust encryption with minimal performance degradation. East-west security (within a data

center or system) can be greatly aided by silicon-enhanced root of trust security technologies, along with firewalls and intrusion protection. It is encouraged to encrypt your data at 'rest' in storage with advanced encryption standards (AES), encrypt 'in transit' via transport layer security (TLS), and encrypt 'in use' via secure hypervisor and encrypted memory transactions.

6. Architectural compatibility

Whether to conduct routine/preventive maintenance, or hardware upgrades on servers, the need to move VMs across the data center and clouds is inevitable. As enterprises modernize and transform their infrastructure with hybrid-multi cloud, careful planning to support hot migration of workloads is paramount. The alternative (cold migration) results in service interruption, scheduling of downtime during periods of low usage (likely after-hours), and managerial overhead.

When hot migrating VMs, it is important to ensure that the architectures where you will be running your VMs are compatible. Migrating across dissimilar CPU brands is not supported by virtualization software vendors, and even migrating the same brand but distinct generation of CPUs presents challenges and suboptimal performance. The reason is that Operating Systems Vendors and Applications Vendors routinely update and optimize their software to take advantage of the latest features introduced in each generation of the CPUs. Thus when live migrating a VM that has powered up utilizing that new capability, that requires that the same CPU technologies be available in destination server for migration to flow seamlessly.

For example, for successful migrations without a degradation of performance and potential application failures, it is essential that factors like the number of cores in the processor along with the memory, storage and I/O architectures are compatible. If you need additional features such as encryption, you may also need to ensure this is made clear from the start. If you are planning to migrate applications between public and private cloud, it is important to be clear about the points of architectural compatibility – and their associated costs – when specifying the public cloud architecture.

7. Maximizing utilization by understanding application frequency

One of the primary reasons to migrate VMs and containers is to maximize efficient use of the infrastructure. This means conducting an analysis of the current workloads across the IT estate and how they are best placed as discussed in point 2 above.

Another vehicle for infrastructure efficiency is VM scheduling. Some applications have significant spikes of activity on a monthly or quarterly basis and are relatively quiet – or totally dormant – the rest of the time. For these periodic

applications, such as end of quarter financial analyses or payroll, it is important to understand the capacity they need when they burst, and what they use the rest of the time. That way, you can ensure that you neither under- nor over-provision. It is worth analyzing the performance trends across various VMs, and allocating them to systems accordingly so that you can maintain a steadily high utilization (many enterprises aim at 80-85 percent) while leaving adequate room for performance spikes.

8. Network infrastructure

VM and container migration relies on maximized network performance when data is going north-south (in and out of the local network) rather than east-west (within the data center network). A migration strategy also needs to include a network strategy, including SDN and virtualized network functions to manage latency. However, network function virtualization (NFV) and SDN don't change the fundamental bandwidth of the connection, and ensuring that network capacity is considered as part of migration plans is essential – especially with hot migrations, or cold migrations which involve large data sets. Maximizing the utilization of the network infrastructure is a crucial component of the VM migration strategy. Network capacity should be neither under-utilized nor in danger of becoming a bottleneck. This is a key consideration of data strategy (point 4 above) as well as overall utilization.

9. Optimizing edge compute

Much of the vast quantity of data produced in the IoT world is going to be processed and initially acted on by virtualized applications running on edge devices, which will cleanse and aggregate data locally, only sending data to be stored to the central data center.

For many industries – for example manufacturing, which produces huge amounts of routine data – it makes sense to run compute and logic at the edge node. Using AI and machine learning, the applications can be trained to only send the important data back to the data center, reducing the data deluge and preserving network bandwidth. For those applications it makes sense to containerize them on the edge devices. Often there will be thousands of identical applications running on identical devices. Containerization makes updating and monitoring much easier.

Intel provides the architecture for the hybrid cloud

We have seen in the section above the key architectural considerations which define a successful VM and container migration strategy. This section of the paper outlines Intel's extensive technology innovation and ecosystem enablement to move, process and store virtualized workloads.

Processing virtualized workloads with Intel® Xeon® Scalable processors

The 2nd generation Intel® Xeon® Scalable processor has greater performance across all benchmarks than previous generations. It allows up to 3.5x greater VM density performance compared to a five-year-old server¹, which allows you to do more with your existing infrastructure, supporting more users and services. For example, data center footprint can be reduced by replacing 20 five-year-old servers with six servers based on two-socket Intel Xeon Platinum 8280 processors. This increased VM density can also deliver up to 59 percent TCO savings at similar performance levels compared to a five-year-old server².

Intel has multiple silicon-level optimizations for VM performance. On the latest Intel Xeon Scalable platforms these optimizations enhance performance and efficiency, helping reduce the virtualization overhead, or tax on the processor. With every new generation of Intel Xeon Scalable processor, the performance benefits of these optimizations scales proportionally with the general processor performance. The performance improvements of the latest Intel Xeon Scalable processors can also reduce the relative burden of VM orchestration (another element of the virtualization tax), ensuring more compute resources are available for application

Intel® Optane™ DC persistent memory increases VM density, TCO and uptime

The expanded platform memory capacity through Intel Optane DC persistent memory facilitates a game-changing increase in the number and density of VMs which can run on each system. For example, compared to a DRAM only-based system, that same system with Intel Optane DC persistent memory supports up to 1.4x more VMware View Planner⁴ VMs on VDI Content Delivery Networks, and up to 8x more VMs when running Redis^{*}, with 8x more capacity⁵. This technology benefits all infrastructures, regardless of placement, enabling more users to be supported within your environment and creating opportunities to enhance and develop service delivery.

Because Intel Optane DC persistent memory is both persistent and very fast, the reboot times for VMs can be reduced.

Intel Optane DC persistent memory also helps increase data security in migratory environments, because the persistency reduces the potential for data loss from system or network failure.

More information about Intel Optane DC persistent memory and virtualization can be found [here](#).

performance. Likewise, the increased processor frequency directly improves the speed and performance of SDN encryption. When it comes to VM migration, moving VMs between Intel technology-based clouds creates zero downtime, compared to up to 18 minutes when migrating from a 60GB common enterprise VM from a non-Intel technology-based environment³.

Intel® Infrastructure Management Technologies offer solutions such as Virtual Machine Device Queues (VMDQ), a component of Intel® Virtualization Technology (Intel® VT) for Connectivity (Intel® VT-c), which optimizes the processing of VM data traffic to improve CPU utilization and bandwidth. Enhanced Intel® Resource Director Technology brings new levels of visibility and control over how shared resources such as last-level cache (LLC) and memory bandwidth are used by applications, VMs and containers. Intel® Ethernet controllers with Application Device Queue (on Intel Ethernet adapters) provide various offloads to enhance SDN/NFV solutions.

Security is enhanced with the latest Intel Xeon Scalable processors with encryption accelerators with Intel® QuickAssist Technology and Intel® Advanced Vector Extensions 512 (Intel® AVX-512). Platform-level security is enabled with Intel® Trusted Execution Technology (Intel® TXT) and Intel® Threat Detection Technology (Intel® TDT).

Finally, Intel® Optane™ DC persistent memory only works with the latest Intel Xeon Scalable processors, taking advantage of advances in the architecture to radically reshape what is possible with system memory performance, density and TCO.

Expanding memory capacity beyond DRAM limits with Intel® Optane™ SSDs

Intel® Optane™ SSDs are proven, as part of a balanced system, to improve container density and performance while reducing cost. Using Intel Optane SSDs and the latest Intel Xeon Scalable processors has delivered a 10x improvement in VMWare vSAN performance, and increased consolidation by up to 30 percent⁶.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks

Performance results are based on testing as of the date set forth in the configurations and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

¹ Up to 3.5x VM density performance: 1-node, 2x Intel® Xeon® Processor E5-2697 v2 on Canoe Pass with 256 GB (16 slots / 16GB / 1600) total memory, ucode 0x42c on RHEL7.6, 3.10.0-957.el7.x86_64, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe, 2*82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, result: VM density=21, test by Intel on 1/15/2019 vs. 1-node, 2x Intel® Xeon® Platinum 8280 Processor on Wolf Pass with 768 GB (24 slots / 32GB / 2666) total memory, ucode 0x2000056 on RHEL7.6, 3.10.0-957.el7.x86_64, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe, 2*82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, result: VM density=74, test by Intel on 1/15/2019.

² Up to 59% TCO savings with Intel® Xeon® Scalable processor compared to 5-year old system. 1-node, 2x Intel® Xeon® Processor E5-2697 v2 on Canon Pass with 256 GB (16 slots / 16GB / 1600) total memory, ucode 0x42c on RHEL7.6, 3.10.0-957.el7.x86_64, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe, 2*82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, result: VM density=21, test by Intel on 1/15 /2019. 1-node, 2x Intel® Xeon® Platinum 8280 Processor on Wolf Pass with 768 GB (24 slots / 32GB / 2666) total memory, ucode 0x2000056 on RHEL7.6, 3.10.0-957.el7.x86_64, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe, 2*82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, result: VM density=74, test by Intel on 1/15/2019.

Improving time-to-deployment with Intel® Select Solutions

Intel® Select Solutions are pre-configured combinations of hardware and software which dramatically reduce the time required for testing and implementation. The Intel® Select Solution for VMware vSAN* provides optimized balance and price performance for VMware vSAN across compute, storage, and networking components. This helps support greater efficiency and performance across SDN and HCI environments.

Conclusion

The ability to migrate workloads to their most efficient location is key to delivering the promise of cloud to increase efficiency and manage system resources optimally. However, to do so successfully requires thought, planning and the right infrastructure. The steps in this paper, along with the technologies described, will allow IT organizations to increase both their flexibility and resource utilization.

Further reading

- **White Paper:**
[From the Data Center to the Cloud to the Intelligent Edge, the Network of Business Runs on Intel](#)
- **Video:**
[Virtuozzo](#)
- **Principled Technologies Report:**
[Migrate VMs from legacy servers](#)
- **Solution Brief:**
[Intel® Select Solutions for VMware vSAN](#)

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction. Example based on estimates as of March 2019 of equivalent rack performance over 4-year operation on virtualization workload running VMware vSphere Enterprise Plus on Red Hat Enterprise Linux Server and comparing 20 installed 2-socket servers with Intel® Xeon® processor E5-2697 v2 (formerly “IvyBridge”) at a total cost of \$796,563 [Per server cost \$39.8K: acquisition=13.7K, infrastructure and utility=4. 2K, os & software=12.2K, maintenance=9.7K] vs.. 6 new Intel® Xeon® Platinum 8280 (costs based on Platinum 8180 assumptions) at a total cost of \$325,805 [Per server cost \$54.3K: acquisition=28.9K, infrastructure and utility=3. 5K, os & software=12.2K, maintenance=9.7K]. Assumptions based on <https://xeonprocessoradvisor.intel.com>, assumptions as of Feb 13, 2019.

³Principled Technologies Report: “Reap the full potential of workload mobility within the cloud and data center by using consistent processor architecture”: <http://facts.pt/8zysd88>.

Performance results are based on testing by Principled Technologies as of January 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. No product or component can be absolutely secure. For more complete information visit <http://www.intel.com/benchmarks>.

⁴ Configurations: VMware ViewPlanner: 1-node, 2x Intel® Xeon® Platinum 8280L cpu on Intel reference platform with 768 GB (24 slots / 32GB / 2666) total memory, ucode 0x4000010 on VMware ESXi 6.7.0, 9808462, 4x Intel SSD DC S4600 Series 2 TB, 2x Intel SSD DC P3700 Series 2 TB, 10 GbE Intel X520 NIC, ViewPlanner 4.0 Beta Release with chrome tab pre-run, Local Mode, 5 iterations, 10s think time, custom chrome profile, HT on, Turbo on, score VM=100, test by Intel on 1/23/2019 vs. 1-node, 2x Intel® Xeon® Platinum 8280L cpu on Intel reference platform with 192GB DDR + 1024 GB Intel Optane DCPMM GB (12 slots / 16GB / 2666 DDR + 8 slots / 128GB / 2666 Intel Optane DCPMM) total memory, ucode 0x4000010 on VMware ESXi 6.7.0, 9808462, 4x Intel SSD DC S4600 Series 2 TB, 2x Intel SSD DC P3700 Series 2 TB, 10 GbE Intel X520 NIC, ViewPlanner 4.0 Beta Release with chrome tab pre-run, Local Mode, 5 iterations, 10s think time, custom chrome profile, HT on, Turbo on, score VM=140, test by Intel on 1/23/2019.

⁵ Configurations: 1-node, 2x Intel Xeon Platinum 8276 cpu on Intel reference platform with 768 GB (12 slots / 32GB / 666) total memory, BIOS PLYXCRB1.86B.0573. D10.1901300453 on Fedora-27, 4.20.4-200. fc29.x86_64, 2x40G, Redis 4.0.11, memtier_benchmark- 1.2.12 (80/20 read/write); 1K record size, KVM, 1 /VM, centos-7.0, HT on, Turbo on, test by Intel on 2/22/2019. 1-node, 2x Intel Xeon Platinum 8276 cpu on Intel reference platform with 192 + 6144 GB (12 slots / 16GB / 2666 DDR + 12 slots / 512GB/ 2666 Intel Optane DCPMM) total memory, BIOS PLYXCRB1.86B.0573. D10.1901300453 on Fedora-27, 4.20.4-200. fc29.x86_64, 2x40G, Redis 4.0.11, memtier_benchmark- 1.2.12 (80/20 read/write); 1K record size, KVM, 1 /VM, centos-7.0, Memory mode, HT on, Turbo on, test by Intel on 2/22/2019.

⁶ Source: <https://www.evaluatorgroup.com/document/lab-insight-latest-intel-technologies-power-new-performance-levels-vmware-vsan-2018-update/> Tests by The Evaluator Group. Configuration details available from The Evaluator Group at <https://www.evaluatorgroup.com/document/lab-insight-latest-intel-technologies-power-new-performance-levels-vmware-vsan-2018-update/>. Performance results are based on testing as of 29th October 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No component or product can be absolutely secure.

Solution Provided By:



Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice Revision #20110804

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com

Intel, the Intel logo, Xeon, and Optane, are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.