Reap the full potential of workload mobility within the cloud and data center by using consistent processor architecture

You can migrate live VMs between Intel processor-based servers but migration in a mixed CPU environment requires downtime and administrative hassle

A study commissioned by Intel Corp.

One of the greatest advantages of adopting a adopting a public, private, or hybrid cloud environment is being able to easily migrate the virtual machines that run your critical business applications—within the data center, across data centers, and between clouds. Routine hardware maintenance, data center expansion, server hardware upgrades, VM consolidation, and other events all require your IT staff to migrate VMs. For years, one powerful tool in your arsenal has been VMware vSphere® vMotion®, which can live migrate VMs from one host to another with zero downtime, provided the servers share the same underlying architecture. The EVC (Enhanced vMotion Compatibility) feature of vMotion makes it possible to live migrate virtual machines even between different generations of CPUs within a given architecture.[1]

Enterprises running hybrid or multi-cloud environments require workload portability to achieve maximum business agility. Two of the key benefits of multi-cloud strategies are flexibility and choice, and the underlying architecture should support these. However, if a customer were to deploy a heterogeneous, or mixed CPU, environment, migrating virtual machines would be cumbersome.

Mixed CPU environments require virtual machines to be shut down before migrating. During this so-called cold migration, users experience some period of downtime before applications resume on the new host. The service disruption often prompts staff to schedule migrations outside of production hours.

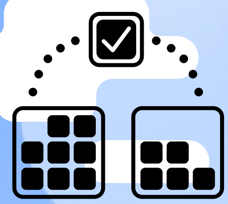

No downtime during live migration

between legacy and current servers powered by Intel® Xeon® processors

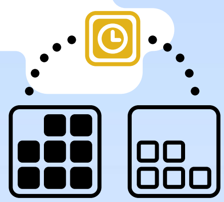

42 seconds of downtime for a cold migration (with shared storage)

Intel processor-powered server to AMD EPYC™ processor-powered server

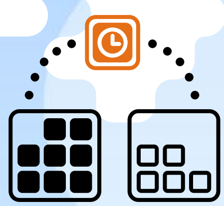

18 minutes of downtime for a cold migration (without shared storage)

from an Intel processor-powered server to an AMD EPYC processor-powered server

CIOs whose data centers have deployed servers powered by Intel Xeon processors and are considering deploying new servers based on the AMD EPYC™ line of processors into this environment should consider how doing so would limit their workload mobility. As AMD notes in their Virtual Machine Migration Guide, when moving VMs between EPYC processor-based servers and servers powered by Intel processors, live migration is not an option.[2]

	Live migration	Cold migration
Intel to Intel	✓	✓
Intel to AMD	✗	✓

In our hands-on testing, we migrated VMs in a homogeneous environment of only Intel Xeon processor-based servers and in a heterogeneous environment consisting of both Intel Xeon processor- and AMD EPYC processor-based servers. We used two new servers powered by Intel Xeon Platinum 8160 processors, one server powered by AMD EPYC 7601 processors, and one legacy server powered by Intel Xeon E5-2680 v2 processors.

We live migrated VMs among the Intel processor-based servers with zero downtime. However, to move a VM from an Intel Xeon processor-based server to one powered by AMD EPYC processors required shutting down the VM for a cold migration.

Because VMware EVC supports migration from Intel processors released as early as 2006,[3] customers deploying Intel Xeon Scalable Processors (also known as Skylake processors) today can be confident that they will reap the full benefits of workload mobility.

About the servers and processors we tested

For our comparison, we chose servers using processors that were roughly comparable in terms of system price. (See the science addendum to this report for complete server specifications.) Each of the servers powered by Intel Xeon Platinum 8160 processors had 48 physical CPU cores and 384 GB of memory. Each of the servers powered by AMD EPYC 7601 processors had 64 physical CPU cores and 512 GB of memory. Our legacy server, powered by Intel Xeon processors E5-2680 v2, had 20 physical cores and 96 GB of memory. Note: Another Principled Technologies study explored the performance of two such servers.[4]

An overview of our test approach

Imagine a data center with an installed base of servers powered by different Intel Xeon processors. The time has come to add new servers, and the team researching the purchase is debating between servers powered by Intel Xeon Scalable processors or servers powered by AMD EPYC processors. To gain insight into the migration options that would be available with each server, we set up the following:

- Two current-generation servers powered by Intel Xeon Platinum 8160 processors
- One current-generation server powered by AMD EPYC 7601 processors
- One legacy server powered by Intel Xeon E5-2680 v2 processors

About VMware vMotion

Three migration scenarios are possible with VMware vMotion:

- Move compute only (shared storage)
- Move compute and storage (non-shared storage)
- Move storage only (not tested; no processor migration)

We tested the first and second scenarios. To test moving compute only, we hosted the VM on external NFS storage presented to both hosts, so that no storage transfer was necessary. To test moving compute and storage, we hosted the VM on direct-attached local storage, so that vMotion had to move VM files from one host's storage to another. We did not test the third scenario as it did not involve migrating between processors.

Using VMware vSphere 6.7, we created a virtual machine running Microsoft Windows Server 2016 Datacenter on each server. We installed Microsoft SQL Server 2016 onto the VM and tested availability with a Microsoft SQL Server database workload, which is used by many customers working in data centers.[5]

We then performed the following migration scenarios:

1. Using VMware vMotion live migration to move an active VM between two current-generation servers powered by Intel Xeon Platinum 8160 processors
2. Using VMware vMotion live migration to move an active VM from a legacy server powered by Intel Xeon E5-2680 v2 processors to a current-generation server powered by Intel Xeon Platinum 8160 processors
3. Using VMware vMotion to move a VM from a server powered by Intel Xeon Platinum 8160 processors to a server powered by AMD EPYC 7601 processors

We used a 60GB VM to represent a typical database VM size. In the scenario without shared storage, the migration time depends on the time required to copy to the virtual disk to the new storage location. The migration time would increase for VMs larger than 60GB, and decrease for VMs smaller than 60GB.

Scenario 1: Migrating an active VM between two servers powered by Intel Xeon Platinum 8160 processors

In this phase of our testing, we migrated an active virtual machine from one of the Intel Xeon Platinum 8160 processor-powered servers to the other under two configurations: with and without shared storage. In both cases, the VM was running a database application serving requests to an external client.

Because live migration is supported between Intel Xeon processors, the database connection stayed up and continued servicing requests throughout the migration period. This was the case with both shared and non-shared storage. An end user accessing this database would experience no interruption, and could work through the migration with no indication that the server hosting the database had changed.

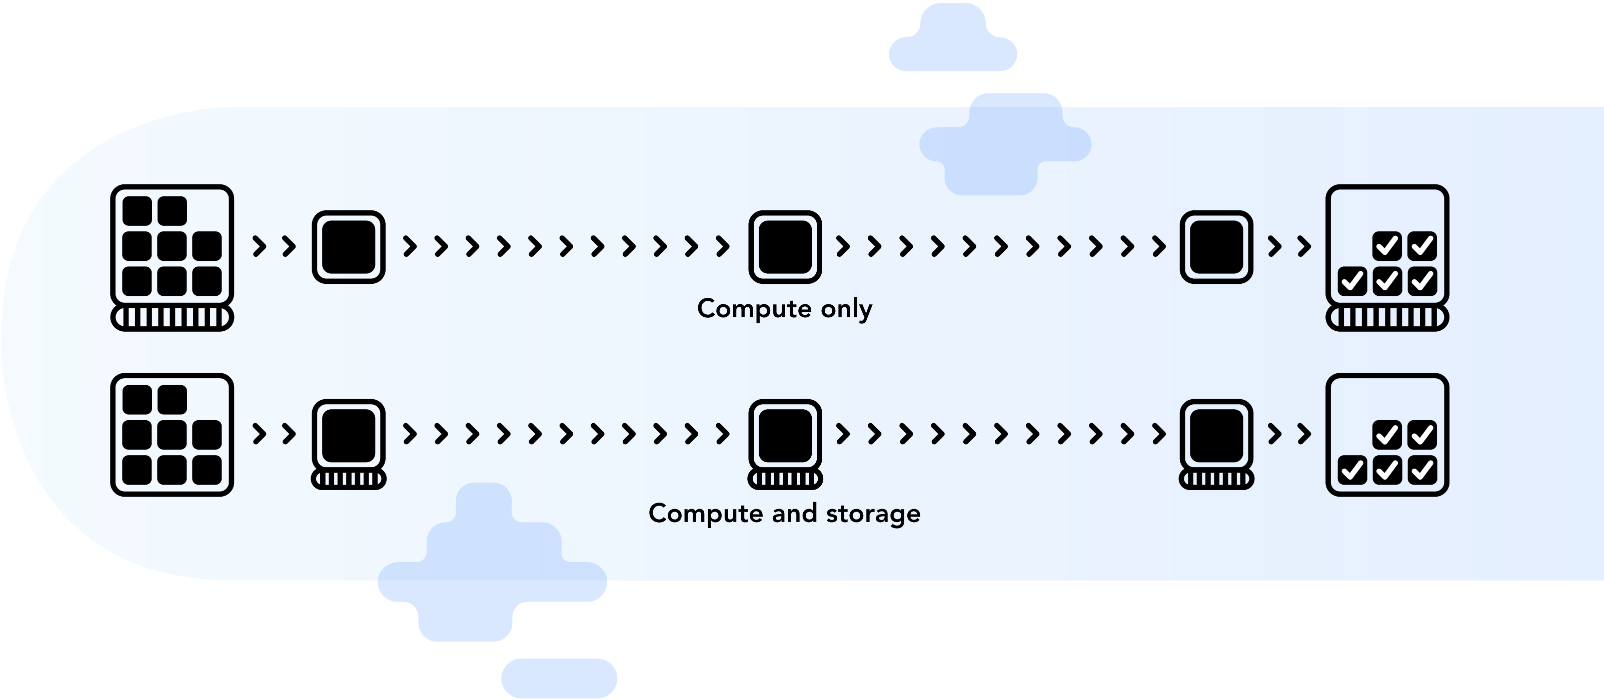

Scenario 2: Migrating an active VM from a legacy server powered by Intel Xeon E5-2680 v2 processors to a current-generation server powered by Intel Xeon Platinum 8160 processors using VMware Enhanced vMotion Compatibility (EVC)

In this phase of our testing, we migrated an active virtual machine running a database application from a legacy Intel Xeon E5-2680 v2 processor-powered server to a current-generation Intel Xeon Platinum 8160 processor-powered server. As in the scenario above, we used an external client machine to send database requests to our active VM and tested one configuration with shared storage and one without shared storage.

We placed both servers in a VMware EVC Cluster and used EVC mode, which allows the servers to match processor features across hosts in the cluster so that VMs can be migrated back and forth smoothly. In our testing, the VM stayed online and achieved a true live migration with zero downtime and no interruption to service.

What the findings of Scenarios 1 and 2 could mean for your company

If your company selected new servers powered by Intel Xeon Scalable processors, you would be able to migrate VMs seamlessly—not only among the new servers but also between legacy servers and new servers. Without live migration, admins would have to schedule downtime for periods that would cause the least disruption for users. For mission-critical VMs, this can quickly turn into a large project requiring resources from technical, managerial, and executive departments, and requiring after-hours work. Instead, that staff can focus their energies on activities that bring value to the company.

Scenario 3: Migrating a VM from a server powered by Intel Xeon Platinum 8160 processors to a server powered by AMD EPYC 7601 processors

In the third phase of our testing, we tried to migrate a virtual machine from an Intel Xeon Platinum 8160 processor-powered server to an AMD EPYC 7601 processor-powered server, both with and without shared storage. We could not perform live migration, unlike in the previous two test scenarios. (Note: VMware Enhanced vMotion Compatibility, which we used in the second phase of testing, is available only for servers powered by processors with the same underlying architecture.[6])

Because live migration is not an option when migrating between a server powered by Intel Xeon processors and one powered by AMD EPYC processors, we had to shut down the active VM running on the Intel Xeon Platinum 8160 processor-powered server before we could migrate it. This is commonly known as a cold migration. In an enterprise environment, this would require users to pause during the migration period. To mitigate the burden this would place on users, IT would likely try to schedule the migration for the time when usage is lowest.

About Intel Xeon Scalable processors

The Intel Xeon Scalable processor series includes many upgrades over previous-generation Intel processors. They have as many as 28 cores (which can enable greater performance and scalability), six memory channels, and up to 1.5 TB of memory per socket.[7]

The series includes four feature configurations: Platinum, Gold, Silver, and Bronze. The Intel Xeon Platinum 8160 processors in the servers we tested each had a 33MB L3 cache, 24 cores, 48 threads, and 3.70 GHz max turbo frequency.

Learn more about Intel Xeon Scalable processors at https://www.intel.com/content/www/us/en/processors/xeon/scalable/xeon-scalable-platform.html

We measured how long it took for VMware vMotion to migrate the virtual machine, both with and without shared storage. With shared storage, from the time we shut down the VM, the database application took 42 seconds to resume servicing requests. This time was dramatically longer without shared storage: 18 minutes, 31 seconds.[8]

Time to migrate a 60GB VM running Microsoft SQL Server 2016
(min:sec)

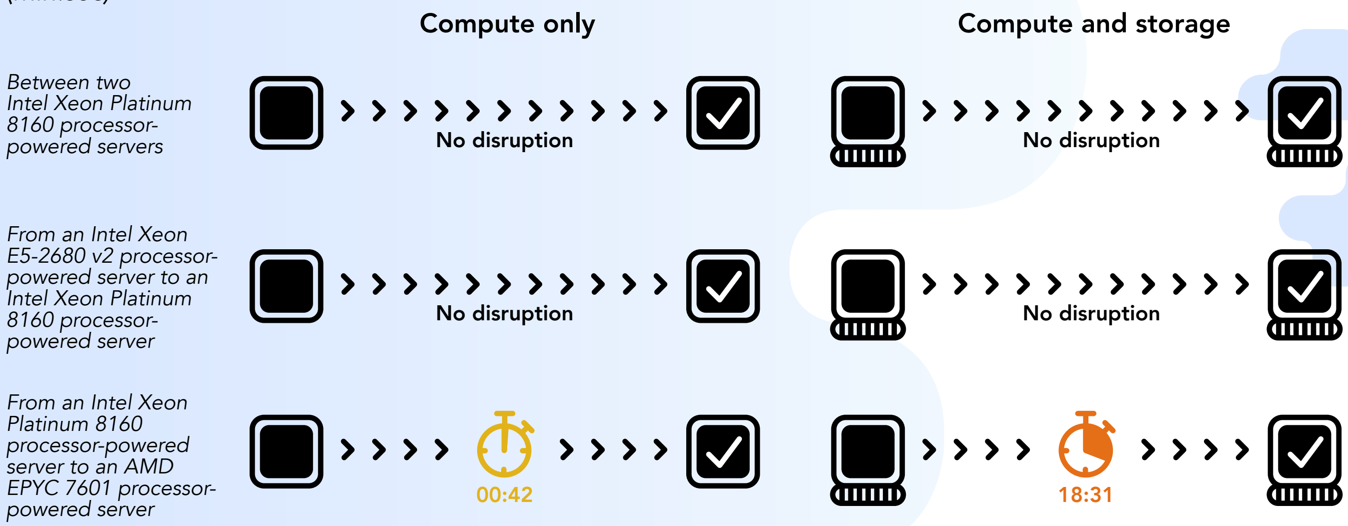

Any length of downtime requires a plan to avoid problems such as losing data from unsaved work or losing customers due to unresponsive services. To avoid these problems, a migration window requires people from multiple teams to work together, often during off-hours, to make sure the downtime does not negatively affect business. The additional time required from team members can quickly become expensive, especially for planning mission-critical VM migrations.

What the findings of Scenario 3 could mean for your company

If your company selected servers powered by the AMD EPYC processor, you would not be able to carry out live migrations between existing servers powered by Intel processors and the AMD processor-based servers. The data center would, in effect, have two separate pools of servers and the mobility of VMs would remain limited for the life of those servers. To avoid service interruptions when migrating VMs, your company would have two choices: (1) performing cold migrations during scheduled downtime, which creates extra managerial overhead and after-hours work, or (2) incurring the extra costs of implementing a highly-available solution."

Conclusion

For years, technologies such as VMware vMotion have made it easy to move virtual machines from one host to another—provided those hosts have the same processor architecture. When you can take advantage of the live migration capabilities of vMotion, your applications stay up and your users stay connected during the migration period. This eliminates downtime and makes scheduling migrations for periods of low usage unnecessary.

As cloud adoption for small and large enterprises has grown, and as companies are increasingly taking advantage of both private and public clouds, planning for workload portability is key. Movement across data centers and clouds is inevitable for modern IT environments and IT experts rely on VMware live migration technology to power their multi-cloud and hybrid cloud strategies.

However, live migration works only between servers with the same processor architecture. We performed migrations with and without shared storage between two current-generation servers powered by Intel Xeon Platinum 8160 processors, from a legacy server powered by Intel Xeon E5-2680 v2 processors to an Intel Xeon Platinum 8160 processor-powered server, and from an Intel Xeon Platinum 8160 processor-powered server to a server powered by AMD EPYC 7601 processors, and confirmed that live migration was an option only in the environment with servers powered by processors from the same vendor.

Between the current and legacy Intel processor-powered servers, the database application running in the virtual machine remained available throughout the migration period and no downtime occurred. When we tried to migrate a VM from an Intel processor-based server to the AMD EPYC 7601 processor-powered server, however, live migration was not possible. Furthermore, the elapsed time until the database connection resumed for our 60GB VM was 42 seconds in the shared storage configuration and 18 minutes, 31 seconds in the direct-attached storage configuration. This downtime can be expensive for businesses to plan around; thus, companies should factor this added expense into their calculations when selecting servers to expand their datacenters on premises or in the cloud.

---

1 EVC and CPU Compatibility FAQ, accessed January 10, 2019, https://kb.vmware.com/s/article/1005764?lang=en_US#q=EVC%20and%20CPU%20Compatibility%20FAQ%20(1005764).

2 AMD Virtual Migration guide, accessed January 10, 2018, https://developer.amd.com/wp-content/resources/VM%20Migration%20Guide.pdf.

3 Enhanced vMotion Compatibility (EVC) processor support (1003212), accessed January 10, 2018, https://kb.vmware.com/s/article/1003212.

4 Servers powered by Intel Xeon Platinum 8160 processors yielded better performance on a mixed cloud workload than those powered by AMD EPYC 7601 processors, accessed April 4, 2019, http://facts.pt/bwnqd4l.

5 The benchmark we used in our testing was DVD Store version 2.1, available from https://github.com/dvdstore/ds21. We downloaded it on December 3, 2018.

6 VMware Knowledge Base EVC and CPU Compatibility FAQ (1005764), accessed December 19, 2018, https://kb.vmware.com/s/article/1005764#Does%20EVC%20allow%20AMD%20and%20Intel%20CPUs%20to%20be%20vMotion%20compatible.

7 Intel Xeon Scalable processors, accessed December 18, 2018, https://www.intel.com/content/www/us/en/processors/xeon/scalable/xeon-scalable-platform.html.

8 We performed three runs to confirm repeatability. The time to resume servicing requests with shared storage ranged from 41 seconds to 45 seconds, with a median time of 42 seconds. The time to resume servicing requests without shared storage ranged from 18 minutes, 20 seconds to 18 minutes, 50 seconds, with a median time of 18 minutes, 31 seconds.

Read the science behind this report at http://facts.pt/op42zx8 ▶

This project was commissioned by Intel Corp.