VMware EVO SDDC

General

Q. What is VMware EVO SDDC?

A. VMware EVO™ SDDC™ is the easiest way to build and run an SDDC private cloud on an integrated system. Based on an elastic, highly scalable, and distributed hyper-converged architecture, compute, storage, and networking resources are delivered from the hypervisor on industry-standard, pre-qualified hardware offered by select partners. EVO SDDC is ideal for enterprises and service providers focused on greater simplicity, faster time-to-value, enhanced security, and lower total cost of ownership (TCO).

Q. What is the workload capacity of a single physical rack?

A. A single physical rack can support up to 1,000 virtual machines (VM) or 2,000 virtual desktops. This assumes average VM of 2 vCPUs and 8 GB Memory, and 160 GB of Storage, and average desktop size of 1-2 vCPUs, 4GB Memory, and 80GB Storage.

Q. How can I buy EVO SDDC?

A. There are two purchasing options for customers: (1) customers purchase the EVO SDDC software suite from VMware and separately purchase the certified hardware from qualified partners or (2) customers purchase both the EVO SDDC software suite and certified hardware from qualified partners. In both cases, customers will receive a fully pre-integrated system with EVO SDDC software pre-installed on the hardware.

Q. Does EVO SDDC work on any hardware?

A. No, to leverage the full capabilities of EVO SDDC requires tight integration with the underlying hardware. As such, hardware and software bill of materials (BOM) have been predefined and are offered by select partner vendors. Please check with your VMware representative for our current list of qualified partners. We will continue building out our partner ecosystem over time.

Q. Is VMware selling and supporting hardware for EVO SDDC?

A. VMware is working with select hardware vendors to offer a partner-branded integrated hardware and software solution powered by VMware EVO SDDC. However, VMware is not in the business of selling or supporting hardware.

Q. For the integrated system, can I mix servers from multiple vendors within a single rack?

A. Within a physical rack, all servers must originate from the same hardware vendor. However, within an EVO SDDC instance with multiple physical racks, users can choose to mix racks from multiple hardware vendors.

Q. Are professional services or training required to operate EVO SDDC?

A. EVO SDDC is intended to be an easy to deploy and manage solution that is fully pre-validated and pre-integrated. Therefore, we believe limited professional services and training are required to operate EVO SDDC. Also, for existing customers already familiar with VMware products, existing skillsets across all products can be leveraged.

Q. Does EVO SDDC replace my existing management tools, such as vCenter Server, vRealize Operations, vRealize Log Insight or vRealize Automation?

A. No, EVO SDDC does not replace your existing management tools. EVO SDDC manages the deployment, configuration, and operations of hardware and VMware software; it does not actively manage VMs and therefore there is no overlap in functionality. Once clusters are up and running, all VM management is through your usual management tools.

Q. Will I be able to install the software myself?

A. No, not at version 1.0. The software will be integrated and pre-installed with qualified hardware BOMs by select partners. We will explore other options, such as professional services or possibly customer installs in the future.

vmware®

Packaging And Pricing

Q. What are the EVO SDDC software components within EVO SDDC?

A. The EVO SDDC software consists of VMware vSphere® Enterprise Plus, vCenter Server™ Standard, vRealize™ Operations™ Enterprise, vRealize Log Insight™, VMware NSX™, and Virtual SAN™. Depending on customer requirements, vRealize Automation Advanced and VMware Horizon® Enterprise are available as optional add-ons.

Q. What is the minimum purchase requirement? How can I add incremental capacity?

A. An initial order per rack requires a minimum of 8 servers. Beyond the initial order, customers can later incrementally add capacity one server at a time.

Q. What is the licensing model for EVO SDDC? Can I get credit for unused software licenses?

A. EVO SDDC is licensed on a per processor perpetual license model. If a customer previously purchased components of EVO SDDC (e.g., vSphere Enterprise Plus or Virtual SAN), he will be able to upgrade to EVO SDDC through the Fair Value Conversion program.

Q. Is EVO SDDC Manager part of vCloud Suite?

A. No, EVO SDDC Manager is not part of vCloud Suite. EVO SDDC Manager is available only in EVO SDDC and is delivered as part of integrated system.

Technical – System

Q. What is the concept of a virtual rack?

A. A virtual rack is an abstraction that allows users to treat compute, storage, and network resources across multiple racks in an EVO SDDC instance as a single pool of capacity (or a virtual rack).

Customers are required only to provide their policy specifications and capacity requirements for workloads. The VMware EVO SDDC Manager™ maps those requirements into steps such as allocating resources, creating clusters, establishing logical network connectivity, etc.

Q. Can there be multiple vCenter Servers and clusters under a single virtual rack?

A. Yes, multiple vCenter Servers and multiple vCenter clusters are supported, but they are subsumed under the concept of a workload domain.

Q. What is a workload domain?

A. A workload domain is a subset of resources carved out of the virtual rack based on specific capacity, performance, and availability characteristics. Each domain can be a cluster and can span across multiple physical racks. Workloads with similar resource requirements can typically be placed in a single workload domain. A single EVO SDDC instance can support multiple workload domains.

Q. How are multiple EVO SDDC-based integrated systems managed?

A. Multiple physical EVO SDDC systems are managed as a single logical pool of resources. There is a single instance of the EVO SDDC Manager that users access to gain visibility into the physical and virtual resources across the racks.

Q. Does Hardware Management Service (HMS) replace hardware baseband control software?

A. Yes, HMS works directly with the management interface of the managed device, thereby making the element manager for the device unnecessary.

Q. What are the power and cooling requirements?

A. Power: Our baseline requirement is 12-14kW per full rack.

Cooling: Air flow of server network ports and switch ports must be in the same direction.

Q. What are the scale limitations of an EVO SDDC instance?

A. The current validation is based on a scale of 8 racks–the main consideration being the spine switch capacity. The software itself can support a much larger number of physical racks.

Q. Is there flexibility in the hardware configurations (e.g., number of servers per rack, memory, storage size)?

A. Within a rack, server configurations are prescriptively defined to avoid unnecessary management complexity. However, we do provide customers flexibility in terms of number of CPU cores, amount of memory, and size of the storage amongst others. Also, across multiple racks, customers will have the ability to select different hardware vendors per rack.

Q. What are the typical hardware configurations?

A. In general, each rack-mounted compute host is a single server with 2 sockets, each with at least 10 cores. The individual components within the server, such as the storage controller and the NIC, are included in the VMware Compatibility Guide (VCG).

Q. How do the EVO SDDC Managers in individual racks communicate with each other?

A. Individual EVO SDDC Managers communicate via the top-of-rack (ToR) and spine switches located on each rack. For further details, please refer to the EVO SDDC documentation.

Q. Will EVO SDDC support bare-metal workloads and containers?

A. At this time, there are no plans to support bare-metal workloads. Containers are specific to workload deployment and can work transparently running inside a VM managed via EVO SDDC. In future releases, we will deliver a full integration of container deployment and orchestration capability. .

Technical – Lifecycle Management

Q. What is the Lifecycle Management (LCM) capability in EVO SDDC? Will the HMS plugin interface be stable over time? How will updates be handled with respect to plug-ins?

A. EVO SDDC has full inventory of all component software and hardware within the system—down to specific editions, versions, and latest patch bundles. On a periodic or as-needed basis, patches and upgrades of all EVO SDDC software components will be released.

The HMS plug-in may require upgrades due to expanded hardware functionality or vendor-initiated fixes. However, every effort will be made to maintain backward compatibility of HMS with any of its upgrades.

The patching and upgrades of the software stack will be based on a dependency and compatibility matrix, including application of concurrent upgrades.

Q. What software bundles are covered as part of EVO SDDC automated Lifecycle Management (LCM) capabilities?

A. EVO SDDC updates and upgrades for the following software bundles are included:

- VMware software bundles, including vSphere and NSX
- EVO SDDC bundles, including EVO SDDC Manager and HMS updates
- Network ToR and spine switch software update bundles
- Low-level software bundles, including firmware, BIOS, and drivers

Q. How long after a component software release will EVO SDDC include the new VMware software?

A. EVO SDDC is part of an integrated software offering that involves interoperability among different components. Currently, it is expected that new VMware software will be incorporated within approximately a quarter after its general availability.

However, emergency updates will be available sooner as required.

Q. Can upgrade and patch bundles be applied to parts of the infrastructure?

A. Yes, upgrade and patch bundles can be applied to specific workload domains or to the entire infrastructure, as determined by the user. Also, at any given point, users have the option to rollback any upgrades or patches if needed.

Q. How is the hardware lifecycle of an EVO SDDC-based integrated system managed? How are new versions of EVO SDDC-based systems handled as they are deployed?

A. A single rack within an EVO SDDC-based system is composed of a prescriptive set of physical devices that are homogenous. As new hardware becomes available (e.g., new generation of CPUs), new rack configurations will also be made available. These new configurations can join older generations of integrated systems. The retirement of older generations of racks will follow industry-standard practices.

Q. What hardware maintenance options are available (e.g., proactive replacement of components or in the case of failure)? How can exchange/replacement of hardware components be handled?

A. Servers and switches can be replaced or upgraded in the field apart from the software/firmware patching and upgrading that is part of lifecycle management. The components used for replacement must be pre-qualified and pre-tested for use with EVO SDDC.

Q. What are some of the options for backup/restore of the management stack and the workloads? Are there any special considerations for disaster recovery with EVO SDDC?

A. The EVO SDDC management stack (EVO SDDC Manager and the workload configuration) can be backed up and restored independent of the workload (consumers of the EVO SDDC software). Some examples of the items that would need backup/restore are:

• Configuration of the management stack itself
• Configuration information of the consumer stack within the management stack (i.e. part of the state maintained by the management stack)
• Information about cloud services in portal (e.g. who has subscribed to what, consumed services account information, etc.)

Through a clustered load sharing approach, the EVO SDDC Manager can maintain and recover its own state in the event of rack or host/switch failures. The management stack runs in high availability mode to enable recovery from network partitions, loss of an EVO SDDC Manager, etc.

Backup/recovery EVO SDDC is not a part of the EVO SDDC software itself. However, partners can provide out-of-the-box integration packs for backup/restore and other value-added services like disaster recovery.

Q. What are the security considerations for EVO SDDC? How do we plan to use NSX security for both workloads and the EVO SDDC Manager? How does micro-segmentation apply to EVO SDDC?

A. Security considerations for workloads and protecting these workloads via micro-segmentation are an inherent part of the EVO SDDC deployment in an enterprise. Using access control lists, not exposing private IP addresses outside of the EVO SDDC environment, and firewalling of VMs per security policy are part of the EVO SDDC deployment. There is no special micro-segmentation/firewalling of service VMs like the EVO SDDC Manager.

Technical – Networking

Q. How is the ToR on the first rack connected to the data center (DC) network? What happens if this ToR fails? How do the other racks cannot connect to the DC network?

A. The ToR connects to the DC aggregation switch using two or more 40Gbps links. It also runs a Layer 3 protocol like OSPF or BGP to connect to the DC network thus providing reachability between the EVO SDDC complex and the DC network. If this switch fails, a second ToR on the first rack can take over. In addition, features like ECMP are used for spraying traffic across links for path redundancy and load balancing.

Q. What is the deployment model for the ToR to spine links? L2 or L3? Is Spanning Tree used?

A. In an EVO SDDC-based system, all the physical racks are on the same L2 network. In the L2 scenario, the ToR switches are dual connected to two spine switches through 40Gbps links, and Multi Chassis LAG (M-LAG) is used. There is no Spanning Tree Protocol used in this scenario – looping is avoided through M-LAG.

Q. Does the first rack have special networking semantics? If so, what are they?

A. The first rack does have some special semantics. For example, the first rack's HMS is the software used to configure and manage the spine switch complex. In addition, this is the rack where a ToR switch connects to the DC network for external connectivity. To realize this, it uses a Layer 2 protocol or Layer 3 routing protocol (e.g., OSPF or BGP) to the DC network.

Q. Can EVO SDDC-based servers be integrated into an existing ToR and spine switch infrastructure?

A. The model for the EVO SDDC-based deployment is "preconfigure and drop-in," where the EVO SDDC-based integrated system is built step-by-step without the need to integrate into existing deployments or handle complicated configurations.

Q. Does EVO SDDC require VXLAN gateway functionality on the ToR switches?

A. There is no need for bridging traffic between VXLAN (used between vSphere hosts) and VLAN (used for communicating between physical/bare-metal hosts).

Q. Is the Inter-ToR connectivity Layer 2 or Layer 3? If it is Layer 2, is there a plan to move to Layer 3?

A. Currently, the ToR and the spine switches are on Layer 2. MLAG is used between the switches for link load balancing. This model also ensures that there is no special configuration (like STP or TRILL) at the links between the switches and also avoids the need for specialized configuration for multicast protocols like PIM at the switches.

Q. Can a deployment avoid a spine switch if it never expands beyond two racks?

A. Due to the prescriptive nature of the architecture, we require a spine switch even with two racks. There are no plans to have a back-to-back ToR connection.

Q. How is NSX set up and configured? Where does the NSX Manager sit? Is this dependent upon vCenter?

A. For details, please refer to the EVO SDDC documentation.

EVO Family

Q. Are there any technical overlaps between EVO:RAIL and EVO SDDC-based integrated systems?

A. Both EVO:RAIL™ and EVO SDDC are based on the foundational VMware hyper-converged components, vSphere and Virtual SAN.

Q. What are the primary differences between EVO:RAIL and EVO SDDC?

A. EVO:RAIL is a compute virtualization solution for vSphere environments, while EVO SDDC enables a data center-scale private cloud based on SDDC architecture. EVO SDDC can span across one or multiple racks with an end-to-end network infrastructure consisting of a leaf-spine topology to carry East-West traffic.

The EVO:RAIL Engine provides a user interface and other "under the hood" functionality for EVO:RAIL. In EVO SDDC-based integrated systems, the EVO SDDC Manager serves as the single interface for managing the logical and physical infrastructure of the private cloud. From this console, the cloud administrator can provision new private cloud resources, monitor changes to the physical and logical infrastructure, manage lifecycle and other operational activities. It performs the following functions:

- Automates configuration and bring up of both underlying hardware (servers/switches) and virtual infrastructure;
- For NSX, it deploys and configures logical networking services such as logical switches, logical routers, logical firewalls, logical load balancers, logical VPN, and distributed security;
- Automates lifecycle management of hardware and software, including patching and upgrade; and
- Provides a new abstraction, workload domains, for creating resource pools across compute, storage, and networking.

Q. At what scale is EVO SDDC-based integrated systems a better choice than EVO:RAIL?

A. While EVO:RAIL can scale beyond 1,000 VMs, we recommend EVO SDDC for environments with 1,000 VMs or greater.

Q. Is there a migration path from EVO:RAIL to EVO SDDC-based integrated systems?

A. We have internally prototyped a migration path from an existing RAIL deployment to an EVO SDDC instance. However, we do not currently have a timeline for releasing automation tools for the transition.

Other

Q. How can I test out the solution?

A. Please contact your VMware sales representative.

Q. Where can I get more information on EVO SDDC?

A. http://www.vmware.com/products/evosddc