

**Go to: Before You Buy or Deploy - Considerations for Design and Procurement**

## Contents

- [1 Introduction](#)
- [2 General Guidelines for SAN/NAS](#)
- [3 SAN/NAS Link Provisioning and High Availability](#)
- [4 Requirements and Best Practices for Storage Array LUNs for Collaboration Applications](#)
- [5 IOPS and other Performance Requirements](#)

## Introduction

This page is intended to provide the requirements when using external storage arrays or Specs-Based.

With Cisco Business Edition or any of the TRCs, the DAS configuration (e.g. number of local drives, types of drives, RAID type) has been designed to provide enough storage performance. Just follow the normal rules described in [http://docwiki.cisco.com/wiki/UC\\_Virtualization\\_Supported\\_Hardware](http://docwiki.cisco.com/wiki/UC_Virtualization_Supported_Hardware) such as CPU, memory, and storage capacity requirements.

## General Guidelines for SAN/NAS

- Adapters for storage access must follow supported hardware rules ([click here](#)).
- Cisco UC apps use a 4 kilobyte block size to determine bandwidth needs.
- Design your deployment in accordance with the UCS High Availability guidelines (see [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/white\\_paper\\_c11-553711.html](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/white_paper_c11-553711.html)).
- 10GbE networks for NFS, FCoE or iSCSI storage access should be configured using [Cisco Platinum Class QoS](#) for the storage traffic.
- Ethernet ports for LAN access and ethernet ports for storage access may be separate or shared. Separate ports may be desired for redundancy purposes. It is the customer's responsibility to ensure external LAN and storage access networks meet [UC app latency, performance and capacity requirements](#).
- In absence of UCS 6100/6200, normal QoS (L3 and L2 marking) can be used starting from the first upstream switch to the storage array.
- With UCS 6100/6200
  - ◆ FC or FCoE: no additional requirements. Automatically handled by Fabric Interconnect switch.
  - ◆ iSCSI or NFS: Follow these best practices:
    - ◇ Use a L2 CoS between the chassis and the upstream switch.
    - ◇ For the storage traffic, recommend a [Platinum class QoS, CoS=5, no drop \(Fiber Channel Equivalent\)](#)
    - ◇ L3 DSCP is optional between the chassis and the first upstream switch.
    - ◇ From the first upstream switch to the storage array, use the normal QoS (L3 and L2 marking). Note that iSCSI or NFS traffic is typically assigned a separate VLAN.
    - ◇ iSCSI or NFS: Ensure that the traffic is prioritized to provide the right IOPS. For a

## UC\_Virtualization\_Storage\_System\_Design\_Requirements

configuration example, see the FlexPod Secure Multi-Tenant (SMT) documentation (<http://www.imaginevirtuallyanything.com/us/>).

- The storage array vendor may have additional best practices as well.
- If disk oversubscription or storage thin provisioning are used, note that UC apps are designed to use 100% of their allocated vDisk, either for UC features (such as Unity Connection message store or Contact Center reporting databases) or critical operations (such as spikes during upgrades, backups or statistics writes). While thin provisioning does not introduce a performance penalty, not having physical disk space available when the app needs it can have the following harmful effects
  - ◆ degrade UC app performance, crash the UC app and/or corrupt the vDisk contents.
  - ◆ lock up all UC VMs on the same LUN in a SAN


## SAN/NAS Link Provisioning and High Availability


Consider the following example to determine the number of physical Fiber Channel (FC) or 10Gig Ethernet links required between your storage array (such as the EMC Clariion CX4 series or NetApp FAS 3000 Series) and SAN switch for example, Nexus or MDS Series SAN Switches), and between your SAN switch and the UCS Fabric Interconnect Switch. This example is presented to give a general idea of the design considerations involved. You should contact your storage vendor to determine the exact requirement.

Assume that the storage array has a total capacity of 28,000 Input/output Operations Per Second (IOPS). Enterprise grade SAN Storage Arrays have at least two service processors (SPs) or controllers for redundancy and load balancing. That means 14,000 IOPS per controller or service processor. With the capacity of 28,000 IOPS, and assuming a 4 KByte block size, we can calculate the throughput per storage array controller as follows:

- $14,000 \text{ I/O per second} * (4000 \text{ Byte block size} * 8) \text{ bits} = 448,000,000 \text{ bits per second}$
- $448,000,000/1024 = 437,500 \text{ Kbits per second}$
- $437,500/1024 = \sim 428 \text{ Mbits per second}$

Adding more overhead, one controller can support a throughput rate of roughly 600 Mbps. Based on this calculation, it is clear that a 4 Gbps FC interface is enough to handle the entire capacity of one Storage Array. Therefore, Cisco recommends putting four FC interfaces between the storage array and storage switch, as shown in the following image, to provide high availability.



 **Note:** Cisco provides storage networking and switching products that are based on industry standards and that work with storage array providers such as EMC, NetApp, and so forth. Virtualized Unified Communications is supported on any storage access and storage array products that are supported by Cisco UCS and VMware. For more details on storage networking, see [http://www.cisco.com/en/US/netsol/ns747/networking\\_solutions\\_sub\\_program\\_home.html](http://www.cisco.com/en/US/netsol/ns747/networking_solutions_sub_program_home.html).

## Requirements and Best Practices for Storage Array LUNs for Collaboration Applications

**The SAN must be compatible with the VMware HCL and compatible with the supported server model used.** A SAN must also meet the following latency storage performance at all time:

- Host-level kernel disk command latency < 4ms (no spikes above) and
- Physical device command latency < 20 ms (no spikes above).

For NFS NAS, guest latency < 24 ms (no spikes above)

There are various ways to design a SAN in order to meet the IOPS requirement for Cisco Collaboration applications (see IO Operations Per Second (IOPS) and therefore to meet the latency storage performance requirements).

The best practices mentioned below are meant only to provide guidelines when deploying a traditional SAN. Data Center storage administrators should carefully consider these best practices and adjust them based on their specific data center network, latency, and high availability requirements.

Other SAN systems such as tiered storage that vary widely by storage vendor could also be used. In all cases, Data Center storage administrators should monitor the storage performance so that the latency storage performance requirements above are met at all time.

## UC\_Virtualization\_Storage\_System\_Design\_Requirements

The storage array Hard Disk Drive (HDD) must be a Fibre Channel (FC) class HDD. These hard drives could vary in size. The current most popular HDD (spindle) sizes are:

- 450 GB, 15K revolutions per minute (RPM) FC HDD
- 300 GB, 15K RPM FC HDD

Both types of HDD provide approximately 180 IOPS. Regardless of the hard drive size used, it is important to try to balance IOPS load and disk space usage.

For Cisco Unified Communications virtual applications, the recommendation is to create a LUN size of between 500 GB and 1.5 TB, depending on the size of the disk and RAID group type used. Also as a recommendation, select the LUN size so that the number of Unified Communications virtual machines per LUN is between 4 and 8. Do not allocate more than eight virtual machines (VMs) per LUN or datastore. The total size of all Virtual Machines (where total size = VM disk + RAM copy) must not exceed 90% of the capacity of a datastore.

LUN filesystem type must be VMFS. Raw Device Mapping (RDM) is not supported.

The following example illustrates an example of these best practices for UC:

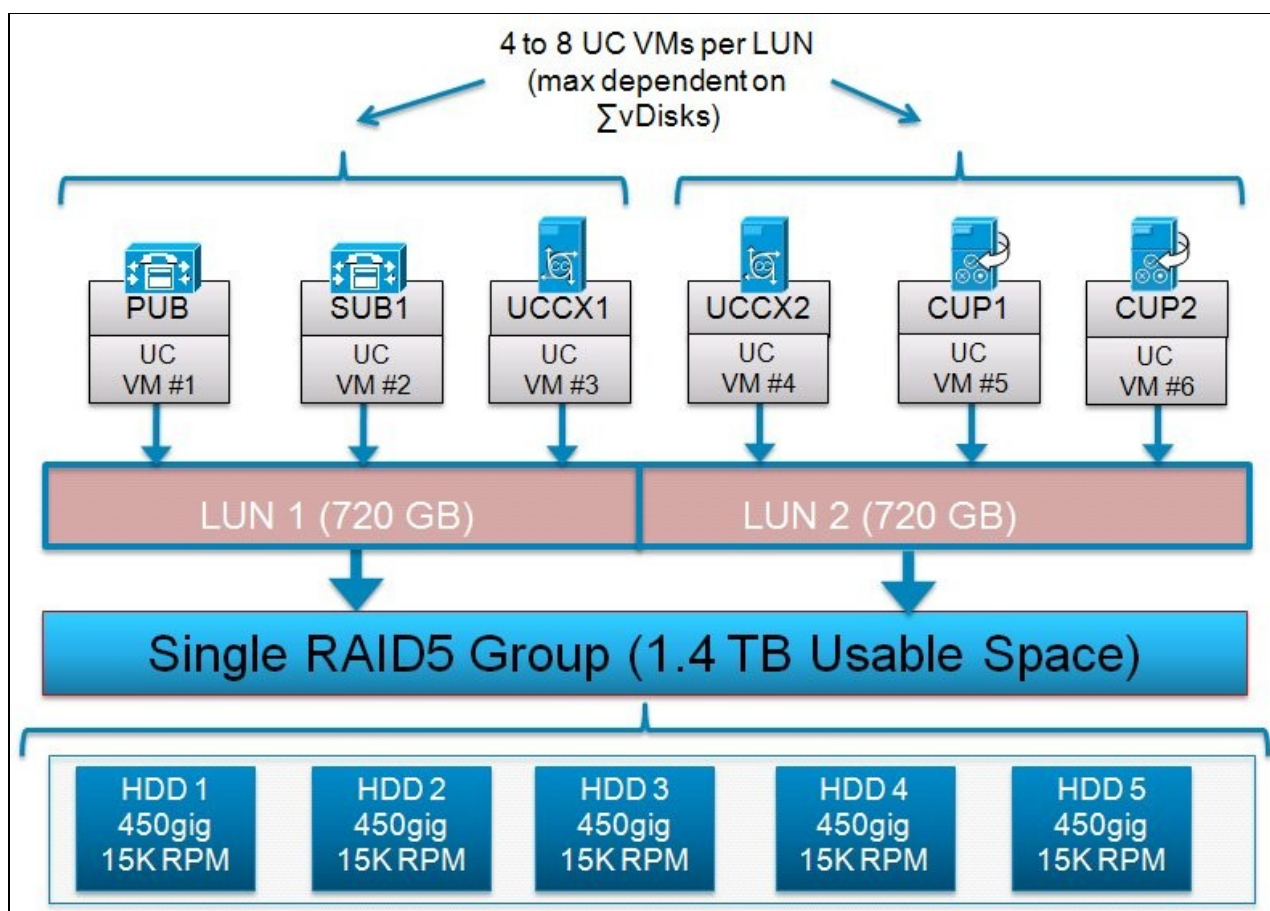
For example, assume RAID5 (4+1) is selected for a storage array containing five 450 GB, 15K RPM drives (HDDs) in a single RAID group. This creates a total RAID5 array size of approximately 1.4 TB usable space. This is lower than the total aggregate disk drive storage space provided by the five 450 GB drives (2.25 TB). This is to be expected because some of the drive space will be used for array creation and almost an entire drive of data will be used for RAID5 striping.

Next, assume two LUNs of approximately 720 GB each are created to store Unified Communications application virtual machines. For this example, between one and three LUNs per RAID group could be created based on need. Creating more than three LUNs per RAID group would violate the previously mentioned recommendation of a LUN size of between 500 GB and 1.5 TB.

A RAID group with RAID 1+0 scheme would also be valid for this example and in fact in some cases could provide better IOPS performance and high availability when compared to a RAID 5 scheme.

The above example of storage array design should be altered based on your specific Unified Communications application IOPS requirements.

Below is a graphic of an example configuration following these best practices guidelines, note there are other designs possible.



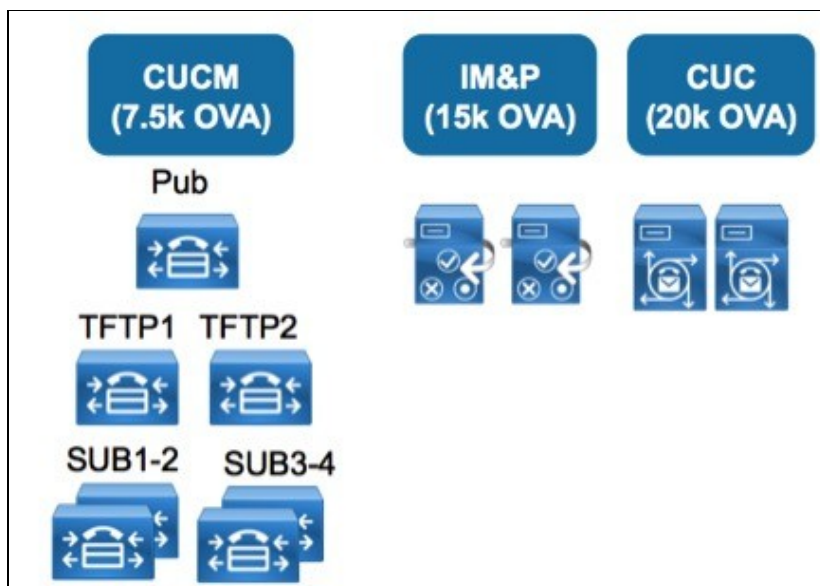
## IOPS and other Performance Requirements

Knowing the IOPS (In/Out Operations per second) utilization of the Cisco Collaboration applications ahead of time will help you design a storage array that will meet the latency storage performance. To find information on the IOPS characterization for a specific Cisco Collaboration application, go the docwiki home page <http://www.cisco.com/go/uc-virtualized>, in the 'At a Glance ? Cisco Virtualization Support?' section, click on that specific Collaboration application, and look for the 'IOPS and Storage System Performance Requirements?' section. IOPS requirements for the storage array is done by using the sum of UC VM OVA IOPS. Note that addressing IOPS requirements may require higher disk/spindle counts, which may result in excess storage capacity. The storage performance should be monitored so that the latency storage performance requirements are met at all time.

### IOPS calculation example

In this example, the deployment includes CUCM, IM & Presence, Unity Connection. Assumptions: 12,000 users/devices, 4 BHCA per user, No CUCM CDR Analysis and Reporting, one application upgrade at a time, one application DRS backup at a time.

Design:



**CUCM IOPS calculation**

PUB and TFTP nodes: Consider the total BHCA for the cluster for the purpose of sizing, 12,000 users x 4 BHCA = 48,000 BHCA. Subscriber nodes: 6,000 users per Subscriber pair, assuming 1:1 redundancy. So 24,000 BHCA per CUCM subscriber pair or an average of 12,000 BHCA per CUCM subscriber. From the IOPS characterization in the CUCM docwiki page, a CUCM node handling between 10k and 25k BHCA produces 50 IOPS. A CUCM node handling between 25k and 50k BHCA produces 100 IOPS. Hence the table below.

	Number of nodes	BHCA per node (for sizing purposes)	IOPS per node	Total IOPS
CUCM Pub and TFTP	3	48,000	100	300 (93%-98% writes)
CUCM subscriber pairs	4	12,000	50	200 (93%-98% writes)

**IM & Presence calculation**

From the IOPS characterization in the IM & Presence docwiki page, IOPS are about 160 when using an OVA with more than 1,000 users.

	Number of nodes	IOPS per node	Total IOPS
IM & Presence nodes	2	160	320

**Unity Connection IOPS calculation**

The IOPS characterization in the Unity Connection docwiki page provides information on the IOPS per node when using the 7 vCPU OVA, refer to the first data column below. We can then calculate the other numbers in the following columns.

IOPS type			

## UC\_Virtualization\_Storage\_System\_Design\_Requirements

	IOPS per node with the 7-vCPU OVA	Additional IOPS per node during peaks	Total IOPS for both nodes	Additional IOPS during peaks for both nodes
Avg Total	202		404	
Peak Total	773	572	1546	1144
Avg Total	10		20	
Peak Read	526	516	1052	1032
Avg Write	192		384	
Peak Write	413	221	826	442

### Total IOPS requirement

The following table shows three types of data: Typical average IOPS during steady state, occasional average IOPS during some operations such as upgrade and/or backup, and additional IOPS during spikes. As you can see, if operations such as upgrades or backups are done while handling calls, the SAN would need to be able to handle more IOPS. It's also a good practice to provide the information to the SAN engineer on additional during peaks. This would allow the SAN engineer to design the SAN cache or increase the SAN performance to handle those peaks. In general, it is a good practice to provide as much information as possible to the SAN engineer, as shown in this table. Again, once the SAN is deployed and with all application running, monitor the SAN performance and ensure the storage latency requirements are met at all time.

	Typical Avg IOPS (steady state)	Occasional Avg IOPS (steady state for a few hours, for example during upgrade or DRS backup)	Additional IOPS during Peaks
CUCM PUB/TFTP	300 (93%-98% seq. writes)	300 (93%-98% seq. writes)	
CUCM call processing Subscribers	200 (93%-98% seq. writes)	200 (93%-98% seq. writes)	
IM&P	320	320	
CUC	404 (95% seq. writes)	404 (95% seq. writes)	Total: 1,144 Read: 1,032, Write: 442
1x DRS		50	
1x Upgrade		1,200 (mostly seq. writes)	
Total	1,224 (~95% seq. writes)	2,474 (mostly seq. writes)	Total: 1,144 Read: 1,032, Write: 442

**Back to: [Unified Communications in a Virtualized Environment](#)**