

THE TOP 10 TIPS & TRICKS TO ROCK YOUR FLASH WORLD

Dell EMC Customer Service has compiled a list of the Top 10 items we believe you should leverage now to optimize your Flash environment. Should you have any questions or if you'd like to add anything to our list based on your own experience using Flash, please visit the [VMAX space on the Dell EMC Community Network](#) and share your thoughts with us.

DEFINING FLASH CAPACITY

- *Allocated Capacity* = What is provisioned to hosts (with virtual provisioning)
- *Physical Used capacity* = Used capacity enhanced by data reduction via compression and/or de-duplication
- *Physical Available Capacity* (aka "writable capacity") = available usable capacity after RAID Overhead, 1024 rounding, hot spare allocation (or capacity for rebuilds)
- *Sizing for capacity* = Physical Used Capacity + Replication overhead + Snapshot needs + Growth factor < *Physical Available Capacity*
- *Logical Capacity* = Max amount of capacity can be provisioned for a given hardware configuration based on overall data reduction. Refer to release notes for each product for details
- Running a storage system at capacity limits can have negative effect during failure scenarios

MANAGE CAPACITY

- In the past, VMAX customers would have to use pre-create specific LUN size and meta volumes. None of that matters any more. Both Unity and XtremIO along with VMAX support simplified storage expansion by simple clicks. It's recommended to provision capacity only as required and grow as needed. No longer do customers have to use old policies and overprovisioning. Expansion of a volume and growing file system requires minimum effort when compared to shrinking a volume or migrating data.
- Delete or unbind the volumes which are not in use by application hosts.
- With Unity systems, unified storage provisioning allows for simple capacity planning – Block, File, and VVols are all supported from the same unified storage pools. Also view current and historical capacity usage and Snapshot storage consumption all from the storage pool level.
- Also with Unity systems, the system can automatically shrink and extend file systems – The system dynamically allocates or frees space from the file system to try to maintain a 70-75% used-to-allocated ratio.
- Using the space reclamation feature in VSI 6.7 or above, you can reclaim unused storage on datastores, hosts, clusters, folders, and storage folders on XtremIO, VNX, VNXe, vVNX, and VMAX. This can be scheduled as needed. It's always recommended to reclaim capacity in guest VM, followed by at datastore using VSI and then at the array pool level as applicable. For database systems start reclaim at the application layer before the others.
- Use EMC Cloud Integrations – CloudArray (VMAX) and CloudPools (Isilon) to augment capacity by seamlessly moving data out of the primary storage system to alleviate capacity constraints

OPTIMIZE PERFORMANCE

- Because the response time with flash is much lower, make sure to use multiple LUNs
- Write rates (VMAX3 utilizes cache) Flash drives have a finite number of write cycles (3K-5K Range). Using *Persistent Cache* feature, available on VMAX3 All-Flash, allows "write buffering", which holds writes in cache longer. Buffering writes in cache greatly increases write hit percentages and also limits the need to

constantly write to the drive. Using persistent cache optimizes write response times and also prolongs flash-media lifespan by reducing write cycles.

- As Flash Drive capacities continue to increase (3.8TB today, 7.6/15.3TB coming in 2H2016) users should be aware of the drive capacity sizes required to deliver the optimal performance and density. Conventional wisdom is that larger capacity flash drives are "better" given the drives can support 1,000's of IOPS. But as capacities double, IO densities are cut in half. As we move to 15.3TB drives, the IO densities become similar to HDDs. The point is users need to be aware that higher capacity flash drives, while capable of sub ms response time, may not be able to meet the IOP requirement for heavy transitional workloads.
- Plan for migration to VVols as it becomes more available with VMware replication and across all platforms.
- Understand your flash platform's load balancing features and best practices.
- In order to maximize system performance, eliminate individual bottlenecks by involving all hardware resources. Spread host connections across as many front-end ports as possible, spread Flash drives across back-end buses, and build storage pools with many drives.
- Leverage Quality of Service Manager (QoS) to limit the amount of host IO that is serviced by the storage system
- Plan for IO Density – Many smaller capacity drives provide better performance than few larger drives
- Data access patterns (hot data), Overall Data Reduction, Host connectivity, Data Services and application integrations will determine which flash system to use for a given application
- Factors effecting overall performance – Application buffering, Host Queuing, SAN Egress/Ingress, Array optimization techniques
- Transports (FC/IP) /Protocols (Block/File) /Connectivity (Speed) will have direct correlation to performance
- IO Parallelism
 - # of Hosts, HW versions (PCI Gen2/Gen3), #Interface Cards/Speeds, # Controllers and multiple application threads

SCALE TO MEET DEMAND

- Start with what you need now, as well as your plans for the future. Have this discussion with your account team now
- Start with a smaller system and small number of controllers/bricks and scale out to meet demand. You can even start from a virtual array if you're just getting started.
- Don't always assume fibre-channel. IP-based networks (NAS, iSCSI) may be a simpler, less expensive solution.
- For Unity systems, large capacities, large drives 6TB/7.6 SSD/1.8 SAS, leverage file for all-flash to get simplicity/scale, new scalable filesystem, pay as you grow with VSA to larger VSA to Appliance

WORKLOAD PLANNING

- Know the signs. Once you get to 75% busy, you're going to have issues.
- With VMAX, use the Workload planner tool. In the past, users would add apps and not know until they implement the last app and things crash. Now they can see before they implement the new if it "fits" and if your capacity will be able to handle...
- For Unity systems, designing for simplicity increases the system's flexibility, and leads to higher, more consistent performance. For instance, it is better to standardize on a single capacity drive, rather than provision a system with multiple drive sizes for different purposes. Standardization can decrease the number of required hot spares, make it possible to build larger pools, and enable greater flexibility for future reconfiguration, if needed.

USE ALL THE TOOLS

- Use MiTrend to estimate data reduction and know what to expect.

- Use your Flash system's PoC Toolkit in evaluations and Proof-of-Concept tests to understand the limits of possible solutions.
- Use EMC Grab and AWR to analyze the application and environment. Understand if speeding up storage is likely to improve user experience.
- Use the Performance Dashboard in Unisphere to analyze performance charts for metrics such as system cache, IO, ports, and disks
- EMC tools such as Unisphere for VMAX, Unity and XtremIO Dashboards allows you to visualize the application performance and SLA compliance. EMC ViPR SRM suite can provide the historical performance trends that can be used as baseline to access IO growth over application lifecycle.

HOST-SIDE BEST PRACTICES

- Enable Round-Robin multipathing using PowerPath or vendor-provided native multipath solutions. PowerPath has predefined policies for each array type.
- Adjust HBA I/O Throttle of the Cisco UCS FNIC HBA to 1024 by changing adapter policy for Linux and VMware.
- Tune the HBA queue depths if possible for low latencies (typically 256 or 128).
- Increase the LUN queue-depth settings from 32 to 256.
- Number of consecutive requests from a single VM from 8 to 64 in vSphere (especially with large sized luns >1TB).
- Queuing – Some level of queuing is good to achieve higher IOPs as long as it doesn't cause increase in response times
 - Queue depth for each target is different 1500 (XtremIO / Unity), 4K for VMAX3 and 64K for DSSD
- I/O scheduling controls how I/O operations are submitted to storage. Linux offers various I/O algorithms (also known as "Elevators") to accommodate for different workloads. Elevators mainly help in reducing seek operations. Adjust OS parameters as appropriate to optimize for low latency (e.g., Linux disk scheduler = noop).
- Generally speaking, use more than a single LUN for a performance-critical application. A minimum of 4 is recommended as a best practice.
- Align disk partitions, file system allocation sizes, and application access block sizes along 8KB boundaries.
- Align the VMFS5 data stores at 1MB offset. Data Stores created using vCenter will automatically align.
- RDM devices >2TB should use GPT with 1MB offset. This will improve the IO performance by 8-15%.

INTEGRATION

- Integrate with AppSync (easier to integrate to database file system). XtremIO iCDM is built from 4 different stack layers; Consistent Scale-Out Performance with inline Data Services, XtremIO Virtual Copies, App Integration and Orchestration, Application self-service.
- VVOLS built-in, SCOM, vCenter, REST API, ESI, VSI, Unisphere visibility into VMs, AppSync for app aware backups, tape backup 3-way integration

DATA PROTECTION

- For Unity systems, D@RE, native async file and block replication, native sync block replication, unified snaps, anti-virus, ndmp backup
- Leverage EMC ProtectPoint Data Protection Software, which integrates the primary storage systems like VMax3 and XtremIO with Industry leading protection platform EMC Data Domain for efficient backup of mission critical applications without impacting application server performance
- VMAX, XtremIO, and Unity support virtual copy creation using snap mechanisms which are instantaneous and space efficient. All these flash platforms support cascaded snaps.
- Use all flash for production, and hybrid flash for backup. Data can be automatically protected using the best low-cost, flexible solution. Consider native

replication solutions for replicating between same platform families and RecoverPoint software for replicating between heterogeneous family platforms.

- Data at rest encryption is supported by EMC primary storage.

SIMPLIFY SUPPORT & MANAGEMENT

- Visit EMC Online Support, Product Pages to download the latest troubleshooting guides and software updates and subscribe to EMC Technical Advisories (ETA) and EMC Support Advisories (ESA)
- Integrate EMC Secure Remote Support (ESRS) to allow quick support and proactive maintenance activities.
 - <https://www.emc.com/customer-services/online-support-tools-resources.htm>
- Take advantage of EMC Ecosystem Integration – VSI, AppSync, ESI, and ESA.
- Use MyService360™ to monitor capacity.
- Get the most stable code and newest features by performing a NDU to the latest released version. NDU activities should only be planned after verifying host timeouts are set to 30 sec, RecoverPoint replication sessions are paused, multiple paths are validated and EMC simple support matrix is reviewed.
- Monitor physical storage usage and plan for capacity expansion.
- Use Software Defined Storage tools such as ViPR for standardization and automation.
- If you have Unity systems, leverage CloudIQ and ProactiveAssist (built off of HTML5 Unisphere) which helps to consolidate alerts and seamlessly send data collects back to EMC.