COMPUTER CENTRE, INDIAN INSTITUE OF TECHNOLOGY KANPUR
KANPUR, UP 208016, INDIA

CC/IITK/09/584
Date:6-04-2018

Sealed quotations are invited from the original equipment manufacturer (OEM) NETAPP, with their principal place of business located at Boeing Avenue 300, 1119 PZ Schiphol-Rijk, The Netherlands, or their affiliate companies registered in India or their certified and authorized implementation and support partner in case the OEM or their affiliates do not bid directly in India for upgrade of the mail storage system at the Computer Centre, Indian Institute of Technology Kanpur, Kanpur UP 208 016, India. In case the bidder is not the OEM then the bidder must submit a duly signed authorization from the OEM. Detailed description of the existing system and the technical requirements for the upgrade are also available in the attached sheets:

Terms and Conditions:

1. All quotations must reach the undersigned by 5PM, April 16, 2018.
2. Quotations must be valid for 75 days from bid opening date.
3. Quotations shall be submitted in two parts.
    - Part-1 (Technical) should contain all technical details cum specifications of the upgrade solution.
    - Part-2 (Commercial) should contain the detailed prices of sub components of the items A to E as described in the "Specifications for the Upgrade of the Mail Storage System" placed in the attached sheets. The total price of the system with and without the buyback option (item E) must be clearly indicated in the bid.
4. Quotations must clearly specify deviations, if any, from the requirements specified in the document.
5. Delivery and Installation: maximum 3 months from the date of receipt of purchase order
6. All software provided in the solution should be perpetual. All licensed software for the hardware must be provided upfront and for the entire possible capacity of the system.
7. Dollar bids must be quoted FOB and CIF.
8. The entire solution should be validated and certified by the OEM
9. IIT Kanpur reserves the right to cancel the bid proceedings at its discretion without assigning any reasons to the vendor.
10. The OEM representatives and bidder if any should present their techno-commercial bid at the Conference Room, Computer Centre, IIT Kanpur on 17-04-2018 at 11:00 AM after bid opening

Amey Karkare 6/4/2018
Amey Karkare
Head Computer Centre

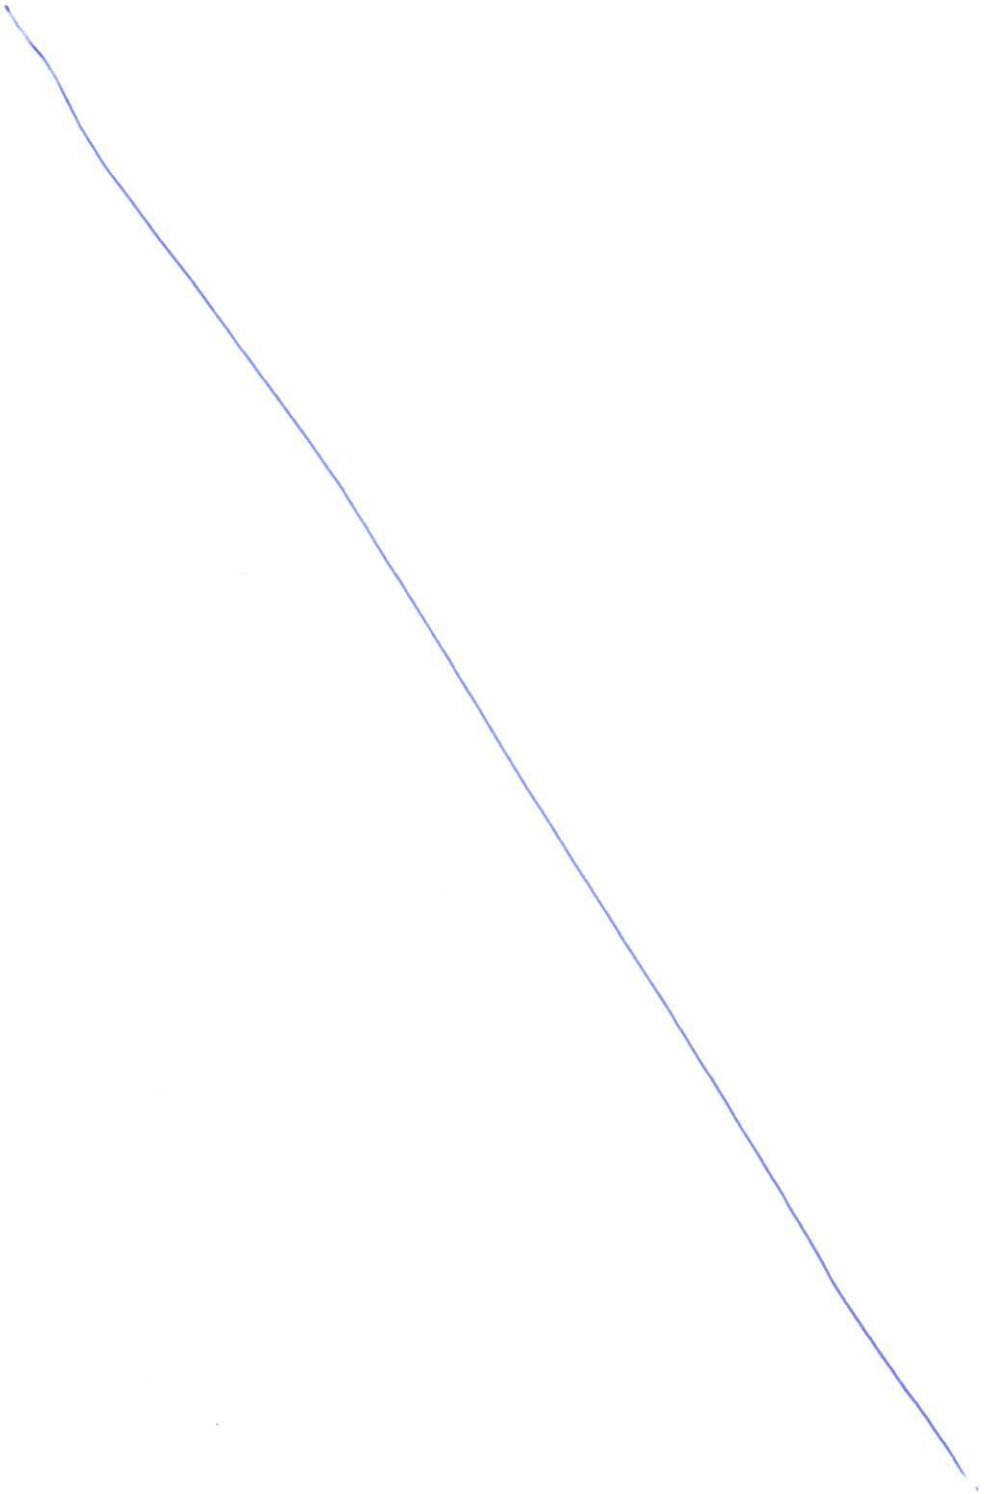

DESCRIPTION OF THE EXISTING SYTEM AND SPECIFICATIONS FOR THE UPGRADE

1. BACKGROUND

The Computer Centre of Indian Institute of Kanpur (CC) has a mail storage system deployed as per the schematic below. The clients access the primary mail storage (dual controller FAS3140) through two NEXUS 5000 series FCOE 10G switches in a fail over high availability mode. There are currently around 12700 users whose data is present in the

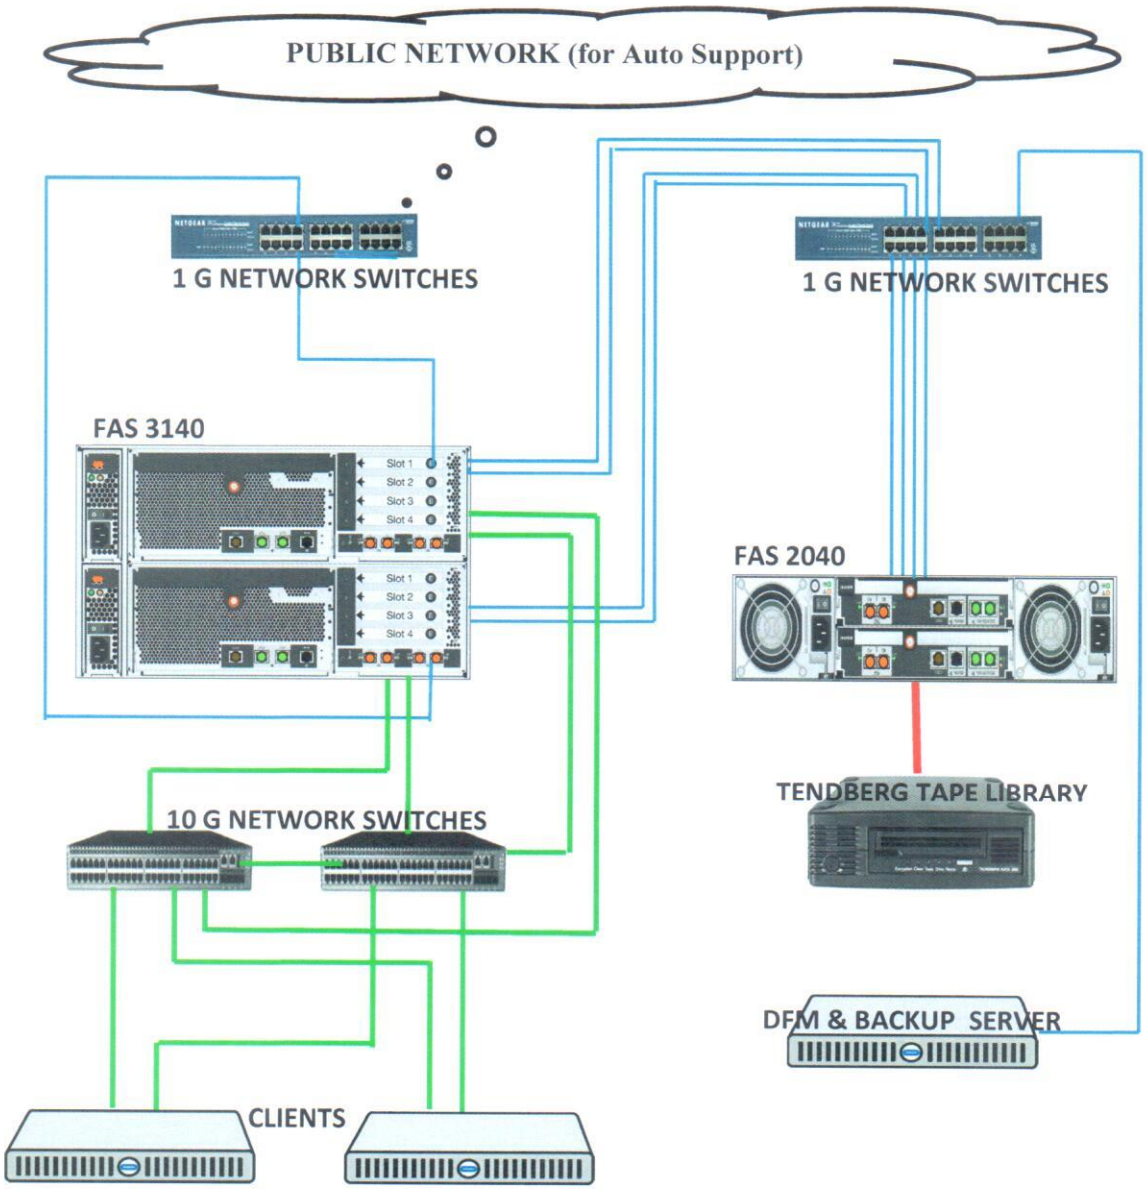

system. The current capacity of the primary storage is 25TB usable on RAID DP. Out of this 20TB is used by Linux based mailing systems and the remaining 5TB is used by the Exchange servers.

The Linux clients mount the home directories via NFS. Five clients are connected to the NEXUS switches via SFP+ DAC (Small Form Pluggable + Directly Attached Copper) cables

(10 cables in all). NFS is served out of 3 "q-trees" from each controller. Thus in all 6 namespaces are exported from the primary storage. The active-active controllers can also failover. Quotas are enabled on the primary storage. These are applied for each user and reset according to system policies each night. The primary storage keeps a total of 48 hourly snapshots and also a total of 7 daily snapshots to back up the data. Currently around 14.5TB of the available space is used. Thus it can restore data of 2 days with a granularity of one hour and the data of a week with the granularity of one day.

A secondary server is connected to this primary storage server by 2 bonded 1G pipes per controller. Backup of the primary data is taken on the secondary through snap-vault. The secondary server can itself be used as a NFS server in case of complete failure of the primary server. However as opposed to the primary server the failover is not automatic but has to be manual. The secondary server also keeps 28 daily snapshots. Thus it can restore data that can be as much as 28 days old with the granularity of one day.

Both the FAS3140 and the FAS2040 are end of support.

A Tandberg Tape Library (LTO 4) is directly connected to the secondary server. The most recent data of the secondary is moved at the beginning of every month to the tape within a window of 36 hours. The tape has one year worth of backup.

There is an additional DFM and Backup Server that controls various activities of the storage controllers. It periodically receives a password file from a designate server and processes it to create home directories of new users. It also syncs the passwords supplied by the designate server to run a NIS server that provides its services to the controllers only. It also computes user quotas according to the policies and resets these quotas every night. Interfaces provided by the controller are used to temporarily change quotas. The controllers also provide monitoring and management services that are made available to the DFM server as well. All servers have RLM/ILO ports for servicing. Auto support is also triggered through the public network.

The windows system takes data through ISCSI and uses snap drive, snap manager for exchange and single mail box recovery to serve users in a mechanism similar to above. Snap mirroring is used to transfer data from primary to the secondary storage.

The Computer Centre now wishes to replace this storage system with another storage system with an increase both in terms of capacity and performance. The rest of the functionality should however remain the same. The primary storage should be replaced by an all flash storage delivering at least 2,50,000 IOPS with 100TB of capacity. The secondary server should now be redesigned and replaced by a storage that can do disk to disk backup for the primary server. At least one year worth of backup should be kept in the disks of the secondary storage. The primary storage would still keep 2 days of hourly backups and 7 daily backups as before for quick restore of data for small periods. In case of failure of the primary server the secondary server could serve as a temporary service provider maybe with some reduced performance.

Amey Datar

The windows aspect to the storage would currently not be required but the options for providing such support should be inbuilt into the storage. Migration issues should be addressed and the restore of the current backup data should be possible in the new solution. The maintenance of the hardware health of the FAS servers/ switches and hardware should also be done during migration. The specifications of the new storage system required are placed in below.

2. SPECIFICATIONS FOR THE UPGRADE OF THE MAIL STORAGE SYSTEM

A. UPGRADE OF PRIMARY ALL FLASH STORAGE:

1. The proposed All Flash storage should be supplied with 100TB usable capacity on SSD (using higher than 960 GB SSD and lower than 8 TB SSD) excluding any RAID and controller OS overhead. Additional 15% usable capacity should be provisioned for snapshots and clones.
2. The proposed system should provide at least 250000 IOPs with installed capacity having workload characteristic of 60:40RW ratio, 8KB block size and not exceeding 1ms latency and should be scalable to 475000 IOPs within the same controller. Benchmarks reports with the latest version of IO meter should be submitted as part of the bid (OEM Sizing documents need to be submitted to support the performance outcome). Bidder should demonstrate the same with latest version of IO meter at the time of UAT.
3. The proposed storage should have active-active dual controllers in highly available configuration with 512 GB usable cache per controller / node. The storage system should be scalable to at least 4 controllers in scale-out architecture.
4. The Proposed storage should support block level inline data de-duplication and inline compression without loss of protection for both file and block data types. In case storage OEM, does not support in-line de-duplication and in-line compression for both block and file data types, an additional 25% capacity needs to be supplied.
5. The Proposed storage should support group and user quotas at various levels of the file system.
6. The proposed storage should be a unified storage system supplied with FC, iSCSI, NFS V3/V4, CIFS/SMB protocols.
7. The proposed system should support single parity, dual parity, mirroring and triple-parity disk protection.
8. The proposed storage should have 4 x 40 or 12*10 Gbps Ethernet ports for NAS and iSCSI connectivity and 4 x 32 or 8 *16 Gbps ports for FC connectivity. The proposed storage should be configured with additional redundant Ethernet ports across controllers for replication. The storage should also be configured with two additional ports for backup if required.
9. The proposed storage should allow monitoring of system performance and capacity utilization. It should also be able to generate reports on the performance and have easy to use GUI based and web enabled administration interface for configuration, storage management and performance analysis tools.
10. The proposed storage should have a single storage OS across SAN and NAS to reduce overall complexity and ensure ease of manageability, upgradability and scalability of the system.
11. The Proposed storage system should have minimum 8 number of 12 Gbps SAS ports across the controllers for backend connectivity.
12. The proposed storage shall have the ability to expand and shrink LUNS / Volumes on the storage online and instantly.

Anuj Tatare

13. The proposed storage should be supplied with thin-provisioning, Quality of services, Snapshots, Clones, Storage Replication, Application aware snapshots and clones (Oracle, SQL, Exchange, VMware and Hyper-V). The license required for the same shall be supplied for the maximum capacity of the offered on the storage model.
14. The proposed storage shall support logical partitioning of controllers in future such that each partition appears as a separate storage.
15. The proposed Storage should have redundant hot swappable components like controllers, disks, power supplies, fans etc.
16. The proposed storage should have the requisite licenses to create point-in-time snapshots. The storage should support minimum 250 snapshots per volume/LUN. The license proposed should be for the complete supported capacity of the system.
17. The proposed storage should support data replication capability across geographically separated locations.
18. The proposed storage should be provided with ability to instantly create clones of active data, with near zero performance impact.
19. The proposed storage should support for industry-leading Operating System platforms including: LINUX, Microsoft Windows, HP-UX, VMware etc. It shall support connecting hosts over iSCSI /NFS / CIFS (SMB) /FC and shall be supplied with any Multipath software, if required, with the solution.
20. The proposed storage should support user and group quotas.
21. Any workstation/server and all software required to install and manage the storage with any required network connectivity should be provided on a preferred.
22. All licenses should be provided upfront

B. UPGRADE OF THE BACKUP SOLUTION:

1. The proposed solution should include D2D backup and restore solution with minimum 250 TB usable with dual disk protection on NL - SAS 7.2Krpm drives. The D2D backup and restore solution should support "Incremental forever" backup at least for a year. Any additional capacity to cater for daily backup copies (@0.25% daily rate of change for 50TB source backup) should be incorporated over and above 250TB capacity.
2. The proposed backup solution needs to be sized to complete all backups in within a 6 hour window.
3. The proposed Backup solution needs to be sized and implemented to retain daily backups for one year on Disks
4. The backup software should be licensed for entire capacity that the proposed primary storage system supports
5. The system should be supplied with necessary software licenses for backup and restore. The proposed backup solution should support Scheduling, Compression & De-duplication.
6. The system should be supplied with necessary licenses for application aware backup for Microsoft Exchange, MySQL, SAP, Oracle, VMware and Microsoft Windows Server 2008/2012
7. The proposed backup storage should be configured with min 2*10Gbps upgradable to 4 Ports. The solution should provide connectivity between primary and backup storage. All accessories and cables for the same should be provided.

8. Any server/workstation and all software required for the backup solution with any required network connectivity should be provided.
9. All licenses should be provided upfront.

C. SWITCHES, FRONTEND AND BACKEND CONNECTIVITY:

1. Two failover redundant 40Gbps 19" (width) rack mountable switches with dual inbuilt power supply and redundant fans should be provided.
2. Switches should be Data Centre Class Switches supporting latest DC technologies like Datacentre Bridging and Priority Flow Control.
3. Each 40Gbps switch should have 16 QSFP+, 40 Gbps ports.
4. Each switch should support:
    a. min 4GB of system memory and 8GB SSD drive
    b. buffer memory size of minimum 16MB per switch
    c. non-blocking wire-speed performance on all 1G, 10G and 40G ports
    d. min aggregate switching capacity of 1 Tbps and a minimum aggregate through-put of 900 Mega packets per second to ensure non-blocking wire-speed performance on all ports Switch should support latency less than 500ns
    e. SDN integration with support for OpenFlow protocol
    f. automation software integration like Puppet for automatic configuration of switches
    g. should support: Network Discovery, VLAN Provisioning, SW upgrades, and bulk command push to many switches
5. Each switch should be provided with:
    a. Three cables that achieve 40 Gbps to 4*10 Gbps breakout functionality (5mt). This would achieve connectivity to existing servers and the proposed secondary backup.
    b. Eleven cables that achieve 40 Gbps QSFP+ to 40 Gbps QSFP+ cables (four 3mt and two 1mt and 5 mt) for primary storage connectivity and inter switch connectivity for future requirements.
    c. Two 10Gbps Short Range (SR) fibre transceivers with cables of 30 mt and 2 numbers of 1 Gbps copper transceivers with cables of 5mt (for connecting to the backbone network of the Institute.
    d. One cable of each type should be kept as onsite spare.
6. All transceivers, cables and accessories for connecting between the storage servers and between the clients must be provided.
7. The OEM should look at the current connectivity and propose the solution.
8. Proposal should include 5 years Next Business Day warranty from OEM and 24x7 accesses to TAC.

D. DATA MIGRATION:

1. Data Migration from the existing environment to the new storage should be factored into the proposal.
2. The estimated downtime for the application based phased data migration activity should not exceed more than an hour.
3. The existing quota mechanisms must be preserved by the migration as well.
4. Total data size for migration purpose would be approx. 12 TB.
5. Storage OEM should be responsible for the migration and must provide services to undertake migration of the existing environment to the proposed environment.

Amey Tatare

6. The scope of this migration encompasses but is not limited to storage data migration, quota migration, authentication migration that may include NIS to LDAP migration, mail box migration, mailing server migration and backup storage migration.
7. This activity must be carried out by storage OEM or OEM certified engineers.
8. Detailed statement of work (SOW) of these services need to be furnished as part of the technical bid

E. BUYBACK:

1. Proposal must also include an optional buyback proposal for the existing storage infrastructure viz. the FAS3140 and the FAS2040.
2. The existing infrastructure will only be handed over to the bidder after the new infrastructure is in production and the buyback option is opted for, after which the bidder is responsible for packing and transporting the old storage/s from the site within a month of satisfactorily implementing the upgraded storage.

F. WARRANTY:

1. The solution should include 5 years comprehensive warranty after installation with 4 hours response time at guaranteed 99.9% uptime for the production environment after commissioning on back to back duly signed support agreement/s with OEMs of components for which the bidder is not an OEM. The OEM/s and the bidder if any should strive to provide much less downtime on a best effort basis and sufficient spares like power cables, disks and SFPs should be provided at site to lower the down time. The details of the spares must be provided with the bid.

G. OTHER:

1. Solution should be provided in a 42-U-rack with sufficient mounting hardware and other accessories.
2. The storage OEM / bidder must own the entire older infrastructure and provide maintenance and break-fix support for the same till the new storage solution comes into production.
3. The entire solution including backup should be validated and certified by the OEM
4. A minimum of 2 SSDs and 2 NLSAS disks should be kept as spares on site so that they can be immediately replaced in case of disk failures without waiting for the delivery of the replacement of the failed disks.
5. Training for two persons on the entire solution should be included.
6. All racks, cables, connectors, switches, power supplies, racks and any accessories to realize the entire solution must be provided.
7. Power and rack cooling needed is to be provided.
8. Preferred operating system for administration/ backup servers is LINUX

H. PAYMENT TERMS:

1. Payments shall be made against a Performance Security through a Bank Guarantee in Indian Rupees for the entire warranty period plus sixty days at 10% of the full amount quoted.
2. In case of Dollar quotations the Performance security has to be deposited at the time of acceptance of the Purchase Order.
3. 90% after delivery and 10% after successful installation, deployment and bringing the machines into production.

I. TIMELINES AND PENALTIES:

1. The entire solution would be deployed and brought into production within 3 months of receipt of the purchase order.
2. Penalty at the rate of 1% per week of the entire amount would be payable in case of delay in implementation and final production till the time the solution is finally deployed and signed off.

Amey Tatare.