



Office of the Vice President
and Chief Financial Officer
Procurement Services
<https://procurement.ufl.edu/>

971 Elmore Drive
PO Box 115250
Gainesville, FL 32611-5250
(352) 392-1331 Fax 352-392-8837

December 19th, 2019

ADDENDUM #1 to the University of Florida ITN 20RL-130 Network Attached Storage scheduled to open on **Tuesday January 14th, 2020, at 3pm** at the University of Florida, Elmore Hall Conference Room, Radio Road, Gainesville, Florida.

This addendum shall be considered part of the Contract Documents for the above mentioned ITN 20RL-130 as though it had been issued at the same time and incorporated integrally therewith. Where provisions of the following supplementary data differ from those of the original document, this addendum shall govern and take precedence. All other terms, conditions, and regulations will apply. This addendum consists of:

- Answers to questions asked prior to the deadline of 5pm, December 11th, 2019

Sincerely,

Rob Luetjen
Procurement Agent III

Please acknowledge receipt of Addendum #1 by signing below and returning this addendum with your proposal. Failure to include addendum with your proposal may result in rejection.

Signature

Company Name

Company Address

City/State/Zip

The Foundation for The Gator Nation

An Equal Opportunity Institution

Vendor Questions ITN 20RL-130 Network Attached Storage

1) **Question:** The University of Florida's ITN mentions the total usable capacity for each of the three sites (6.32PB, 6.32PB, and 1.73PB) to support both the SMB and NFS data environment. Could you provide an estimate of the percentage of capacity associated with each environment per site (i.e. 70% SMB and 30% NFS at Active site #1 and Standby site #2, 50% SMB and 50% NFS at Remote site #3).

Answer: See pg. 6, sections 4.a and 4.b for workload protocol split. As for disk capacity, system should be able to work with any ratio of SMB/NFS (99/1, 75/25, 50/50) storage. The standby site should be able to handle the same ratio should the active site fail. The remote site should not have direct client access, so access will be via backup/restore protocols.

2) **Question:** The ITN's storage capacity section (1.2.1 5a) uses the word "logical" in association with "usable capacity". Please confirm that the University of Florida would like 14.37PB of "usable capacity" as defined within the ITN's definition section (5.57).

Answer: logical usable capacity refers to the given definition of usable capacity at a logical not physical level. For example, if the storage vendor utilizes fixed block sizes and the dataset contains lots of small files, then the actual used physical capacity will be higher than logical size. This must be accounted for. Do not include data reduction. Do include requirements such as SSD overprovisioning, RAID parity, hot spares, etc.

3) **Question:** "Section 1.2.1 Required Features #5a" states a requirement for a combined 14.37 PB total usable capacity (**logical** usable capacity).

"Section 5 Definitions - 5.57" defines "usable capacity" as...

The amount of data that can be stored in a system in the absence of any data reduction (thin provisioning, data compaction, data compression, and data deduplication), but including any overhead due to RAID parity and/or SSD over-provisioning.

If the term "logical usable capacity" is intended to depict 14.37PB of effective capacity after data reduction, could you please provide the breakdown of the TYPE of data as a percentage of the requested capacity by location. This will better ensure more accurate sizing based on anticipated data reduction rate. Example...

➤ Active Site #1:

- SMB: Uncompressed File (Word, Excel, etc.) = 40%
- SMB: Compressed File (Images, LIDAR, Video, etc.) = 10%
- SMB: Encrypted Data = 20%

- NFS: Uncompressed File (Word, Excel, etc.) = 10%
- NFS: Compressed File (Images, LIDAR, Video, etc.) = 10%
- NFS: Encrypted Data = 10%
- *any other data type that may potentially skew the anticipated data reduction rate.

Answer: According to the requirements in the ITN, the provided solution must not use data reduction techniques to arrive at the provided capacity numbers.

4) **Question:** Does the University of Florida want to own the storage assets or consider Storage-as-a-Service (pay for capacity utilization on a cost per GB) over the five year term? If capacity utilization is acceptable, how much capacity should be assumed for each year (i.e. yr1 = 6PB, yr2 = 8PB, yr3 = 10PB, yr4 = 12PB, yr5 = 14.37PB) as well as the capacity needed for each protocol SMB, NFS, and Object Storage per year (i.e. yr1 = 4PB SMB, 2PB NFS, 1PB Object).

Answer: We have traditionally used ownership but would consider other models. If another model is used, the ownership model should also be provided. Provided capacity numbers are projections of capacity needed at the time of the award including 3 years of 20% growth. These numbers are calculated based on on-premises use and not being able to use more than 85% of capacity. Reversing these formulas results in logical usable capacity needed at time of award of:

Performance: Yr0=1.03PB (add 20% per year)

Capacity: Yr0=2.07PB (add 20% per year)

Standby: Yr0=3.10PB (add 20% per year)

Remote Site: Yr0=870TB (add 20% per year)

Refer to answer for question 1 for capacity split % for SMB/NFS and calculate accordingly. We don't know how much object storage we will need per year as it is not yet a current offering.

5) **Question:** Would the awardee be able to leverage the University of Florida's existing Microsoft Enterprise Agreement for licenses.

Answer: Any vendor provided solutions must identify necessary licenses to conduct all activities.

With respect to the Tier A data set:

6) **Question:** In your existing solution, how is the 2.1PB of data presented today? Is it in 1 file system/volume, or multiple file systems/volumes? If it is in multiple file systems/volumes, can you provide specifics on how it is organized for example:

- i. Tier_A.Volume_A = 500TB
- ii. Tier_A.Volume_B = 700TB
- iii. Tier_A.Volume_C = 900TB

Answer: See the answers to questions 3 and 9 to understand how the performance and capacity numbers were calculated and provided and the relationship between the current in-use numbers vice the projected 3-years-out from date-of-award numbers. Currently, Tier A is 2 clusters and tier B is 2 clusters. Offsite tier is 1 cluster. Each cluster is a single file system/volume spread across a group of nodes providing a wide arrangement of shares/exports. Each tier is separate, and each user can only view their shares+exports. NFS exports are presented in a single namespace per NAS cluster. SMB shares are presented as either separate UNC paths, or nested UNC paths per NAS cluster, as per user requirement.

7) **Question:** Can you share details about the average file size or the total file count within the 2.1PB of Tier A.?

Answer: The 2.1PB of Tier A is a capacity projection for date of award that includes 20% growth (per year) over three years. Refer to answer to question 3 and question 9 for additional info that may be used to come up with average file size. Total current count exceeds 875 million files and 1.25 billion Logical Inodes (LInes) for Tier A (primary site).

8) **Question:** In the desired solution, are you looking to modify the way the Tier A data is organized? If yes, please provide details about how the organization of the file systems/volumes will change.

Answer: We are not dictating the solution. It is up to the respondents to provide the University with the best technical and cost combinations of storage while meeting the capacity and performance requirements. We included references to HSM to represent our wiliness to explore these options. Review of hot/cold data showed: Archive potential for Tier A: 41.9% by capacity for files not accessed in 1+ years. Refer to answer to question 3 and question 9 for additional info on capacity numbers.

9) **Question:** How are the file systems/volumes used? For example, are they home directories for faculty and students? Are they departmental shares? Are they targets for research data generated by applications, etc.... Any details about the workflows of machine-generated data would be useful for us to understand.

Answer: Yes, we have all these use cases.

10) **Question:** The ITN mentions the need for 20,000 connected clients to Tier A with 2,500 active clients. For clarity, are you saying that 20,000 workstations are always connected with active drive mappings or mounts, with then 2,500 of those connections actually moving data? Or, are you stating that 20,000 clients will have access to the data, but at any one time, the highest number of connections (moving data or not) would be 2,500?

Answer: This is represented in the ITN via the use of peak and median. Connected vs. active client is similarly represented. Numbers stated within the ITN are based on current metrics.

These numbers and percentages were scaled up from captured metrics based upon the requirement that performance grow with capacity and we accounted for capacity growth of 20% per year for three years from the expected date of award.

11) **Question:** If there are multiple file systems/volumes, do all the connections need to interact with all the file systems/volumes or is it split? so for example:

1. Are the connections split like this:
 - a. Tier_A.Volume_A = 1,000 connections
 - b. Tier_A.Volume_B = 1,000 connections
 - c. Tier_A.Volume_C = 500 connections
2. Are the connections common like this:
 - a. Tier_A.Volume_A = 2,500 connections
 - b. Tier_A.Volume_B = 2,500 connections
 - c. Tier_A.Volume_C = 2,500 connections

Answer: Option 1 above is more representative of this approach. The numbers are for the total number of clients connected to the entire Tier and may be split to different shares/exports. How the vendor splits the data or capacity up into volumes or other logical divisions doesn't matter as long as the requirements for per share/export connections on page six (3-O and 3-P) are met and the system and/or tier can provide the performance metrics indicated.

12) **Question:** With respect to the client protocols, SMB and NFS, can you detail how the file systems/volumes of the Tier A data set need to be accessed? We are looking to understand if the file systems/volumes will be SMB only, or NFS only, or mixed SMB & NFS where user and permission mappings are required. For example:

- i. Tier A. Volume_A = SMB Only
- ii. Tier A. Volume_B = NFS Only
- iii. Tier A. Volume_C = SMB & NFS with bi-directional user and permission mapping.

Answer: Example is correct. Most shares/exports will be accessed via a single protocol, but some will be accessed via both concurrently.

With respect to the Tier B data set:

13) **Question:** In your existing solution, how is the 4.22PB of data presented today? Is it in 1 file system/volume, or multiple file systems/volumes? If it is in multiple file systems/volumes can you provide specifics on how it is organized, for example:

- i. Tier_B.Volume_A = 2,000TB
- ii. Tier_B.Volume_B = 2,000TB
- iii. Tier_B.Volume_C = 220TB

Answer: All listed presentation approaches. Each tier is separate, and each user can only view their shares+exports. NFS exports are presented in a single namespace per NAS cluster. SMB shares are presented as either separate UNC paths, or nested UNC paths per NAS cluster, as per user requirement.

14) **Question:** Can you share details about the average file size or the total file count within the 4.22PB of Tier B?

Answer: The 4.22PB of Tier B is a capacity projection for date of award that includes 20% growth (per year) over three years. Refer to answer to question 3 and question 9 for additional info that may be used to come up with average file size.

Total current count exceeds 685 million files and 830 million LINs for Tier B (primary site).

15) **Question:** In the desired solution, are you looking to modify the way the Tier B data is organized? If yes, please provide details about how the organization of the file systems/volumes will change.

Answer: See response to question 7.

16) **Question:** How are the file systems/volumes used? For example, are they home directories for faculty and students? Are they departmental shares? Are they targets for research data generated by applications? Are they archive targets for legacy data? etc... Any details about the workflows of machine-generated data would be useful for us to understand.

Answer: See response to question 8.

17) **Question:** The RFP mentions the need for 7,000 connected clients to Tier B with 550 active clients. For clarity, are you saying that 7,000 workstations are always connected with active drive mappings or mounts, with then 550 of those connections actually moving data? Or, are you stating that 7,000 clients will have access to the data, but at any one time, the highest number of connections (moving data or not) would be 550?

Answer: See response to question 9. Note that this is an ITN and not an RFP.

18) **Question:** If there are multiple file systems/volumes, do all of the connections need to interact with all of the file systems/volumes or is it split? For example:

1. Are the connections split like this:
 - a) Tier_B.Volume_A = 300 connections
 - b) Tier_B.Volume_B = 200 connections
 - c) Tier_B.Volume_C = 50 connections
2. Are the connections common like this:
 - a) Tier_B.Volume_A = 550 connections
 - b) Tier_B.Volume_B = 550 connections
 - c) Tier_B.Volume_C = 550 connections

Answer: See response to question 9.

19) **Question:** With respect to the client protocols, SMB and NFS, can you detail how the file systems/volumes of the Tier A data set need to be accessed? We are looking to understand if the file systems/volumes will be SMB only, or NFS only, or mixed SMB & NFS where user and permission mappings are required. For example:

- i. Tier_B.Volume_A = SMB Only
- ii. Tier_B.Volume_B = NFS Only
- iii. Tier_B.Volume_C = SMB & NFS with bi-directional user and permission mapping.

Answer: See response to question 10

20) **Question:** With respect to the standby and DR sites, under normal operation, where the primary site is fully operational, what, if any, activity will happen at the standby and DR sites?

- a) Will the data just sit and do nothing, just ready for a failure event?
- b) Will the data occasionally be read?
- c) Will the data be both read and write?
- d) Is it fair to assume that under normal circumstances both sites will be un-used or are they in some way used while the primary site is fully functional? If they are in use, can you please detail in what way.

Answer: Our expectation is that the Primary will write to Standby and the standby will write to remote site. Any variation from this should be noted in the proposal.

21) **Question:** Is the 1.73PB of data needed for the “Remote Site” a subset of the data from the primary site “Tier A” + “Tier B” or is this a 3rd data pool?

- a) If it is a subset of Tier A and/or Tier B data how is this subset identified?
- b) If this is a 3rd data pool, not included within Tier A or Tier B, can you please provide details about how this data is organized and used?

Answer: Please see answers for question 3 and 9 to understand relationship between the tiers, capacity projections, and current numbers. Currently our “remote site” is a subset of current capacity of Tier A. We project, expect, intend and have documented the requirements for capacity for “Remote Site” to be a subset of primary site (both Tier A and B). The important aspect is that it must be able to grow to match the end state full capacity of the primary site regardless of the current stated capacity needs.

With respect to performance:

22) **Question:** Are the throughput numbers requested for a single client or for the aggregation of all the clients?

- i. If it is for a single client, can you please specify what type of client this is?

Answer: Throughput numbers listed under page 6, section 4 are aggregated, unless specified otherwise.

23) **Questions:** For Tier A, the RFP requests 4,800,000 file system events per second - can you please define what a file system event is in this context?

Answer: Any discrete command to the API that interacts with the files or filesystem (open, close, read, stat, dele, etc).

24) **Question 1.2 Scope of Work Page 5:** Is a lab system required as part of this proposal? If so, what are the minimum performance and capacity requirements needed for the lab?

Answer: Yes, a separate permanent physical lab system is required. No defined performance and capacity requirements are provided as the intention is for it to function as a testing bed for software/firmware upgrades, configuration changes, design changes, and other as-needed tasks so that these tasks do not interfere with or break the production systems. Lab nodes and hardware should be the same generation hardware as the proposed solution. Some amount of space to perform tests of specific share and/or export configurations and settings is required. At least 20 TB of usable capacity will suffice. High performance isn't required, but it must be similar enough so that results from the necessary tests defined in the ITN will apply to the corresponding production systems.

25) **Question 1.2 Scope of Work Page 5:** The requirements state that a POC test platform must be provided for minimum of 5 working days. Is there a need for a lab system beyond a 30-day maximum?

Answer: The test platform noted in the ITN required during the evaluation periods prior to selection is different from the permanent lab system that is part of the design. This test platform is only needed for the ITN-defined period as part of the evaluation. Once the evaluation and test phase is complete and the final award is made, the POC test platform will no longer be needed. Any vendor provided proof of concepts should be completed within 2 weeks from the start of our systems verifications/validations/testing.

26) **Question 1.2 Scope of Work Page 5:** Lab environment system must be included in the deal - Is one system considered enough? Or are you requiring 2 fully redundant systems so that site replication and site failover can be tested? Or can the 2nd site failover testing be done using a software only license on UF existing server hardware?

Answer: Two separate physical systems may be necessary to complete the full range of configuration and full production replication functionality testing defined in the ITN and intended as part of the final solution design. However, we leave it up to the vendor to make this designation. Production UF systems must not be impacted by the testing.

27) **Question 1.2 Scope of Work Page 5:** SOW for PS is required to be included. Does it require migration services? What is the timeline for the migration of existing data?

What are the maintenance windows that need to be orchestrated around for migration?
Would migration be requested for full migration or assisted support only where PS supports the UF staff and the UF staff does the bulk of the migration work?

Answer: Section 1.2, paragraph 2 states "the complete installation for the replacement project" and that it should "detail all steps." As migration of existing data would normally be considered part of the replacement, it should be listed accordingly. Professional Services (PS) for migration is only required if the migration must be performed by vendor. In the interest of a successful migration outcome, we defer to vendor recommendation on whether to let PS handle the migration entirely or if PS just assists UF staff during the migration period. The time it takes to migrate will depend on the technology employed and the impact on our current operational systems.

28) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting (c):** Does the solution need to integrate with any other 3rd party tools to support the event monitoring and alerts notification requirement under 1.2.1.1.c?

Answer: Yes. SNMP, Syslog, and SMTP are standards-based communications protocols. UFIT will be responsible for configuring the third-party tools, but your solution must support the protocols.

29) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** Where in the directory tree structure on the Isilon cluster(s) are quotas being used today?

Answer: Current Isilon quotas are defined on multiple nested directory levels. This may include any directory components of "/ifs/\$customer/\$share/\${folder|userid}/" or even deeper structures.

30) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** Is the quota requirement feature needed for the active files/directories or for snapshots?

Answer: Both. Refer to 3 Capability number N. "snapshot space use included in volume or quotas."

31) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** Do quotas need to be applied at the user, group, level for a directory or volume?

Answer: Yes.

32) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** Is there a requirement for hard (write prevention) or soft (notification only) quota management?

Answer: Hard quotas are required. Advisory/notification quotas are highly desired.

33) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** How many quotas are there on the Isilon cluster today?

Answer: Tier A has over 25,500 quotas defined; Tier B has over 2,100 quotas defined.

34) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** What is the maximum directory depth for existing quotas?

Answer: We do not restrict directory quota depths. Most existing quotas are less than 5 directories deep.

35) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** Are there any overlapping directory quotas in place on the Isilon cluster?

Answer: Yes. We refer to these as "nested quotas" in Section 3.j.

36) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** How is quota enforcement governed on the Isilon cluster? Is it on by default or an exemption for unique use cases?

Answer: Enabled by default. Every share/export has at least one quota defined as part of our process.

37) **Question 1.2.1 Technical Specifications Page 5 Audit/Monitoring/Reporting:** Is there any quota inheritance in place today where a child is linked to a parent?

Answer: We do have "default-user" quotas enabled on some shares, in where all users share a per-user linked hard threshold, allowing us to change the default without having to modify each user's quota.

38) **Question 1.2.1 Technical Specifications Page 5 Availability/Reliability (b):** Is the RTO of 60 minutes under 1.2.1.2.b required for the entire active site or selective data sets based on criticality or business impact?

Answer: Entire active site (Section 3.f).

39) **Question 1.2.1 Technical Specifications Page 5 Availability/Reliability (b):** Is the RPO of 60 minutes on page 5 required for both Tier A and Tier B?

Answer: Both Tiers A and B.

40) **Question 1.2.1 Technical Specifications Page 5 Availability/Reliability (c):** Is the RTO of 60 minutes on page 5 required for a granular restore or an entire active site reset from the secondary and/or remote site(s)?

Answer: Entire site (all file services provided by Tier A and Tier B at primary site.) The RPO and RTO listed on page 5 under 1.2.1 (required section) 2b and 2c apply to the active and standby sites only. The RPO and RTO for the remote site is covered on page 7 under 4. System Metrics -> c. Remote site -> ii Total local site disaster and recovery (active and standby sites down).

41) **Question 1.2.1 Technical Specifications Page 6 Capability:** There is a minimum requirement for 2,000 snapshots per share/export. What snapshot counts are being realized today on the Isilon cluster?

Answer: Our ITN details should have listed a requirement for 1000 snapshots per share/export minimum (2000 is incorrect). Normally, not more than 60 per share. However, concerns over e-discovery and legal hold requirements may require retaining the snapshots for longer intervals.

42) **Question 1.2.1 Technical Specifications Page 6 Capability:** Is there any snapshot reserve (space allocation) configured today? If so, what percentage is being used?

Answer: The current Isilon clusters do not use snapshot reserves. Snapshot usage are included in the share/export quotas.

43) **Question 1.2.1 Technical Specifications Page 6 Capability:** Are snapshots being kept indefinitely, manually or automatically deleted?

Answer: Not indefinitely. They are kept and maintained on a schedule (deleted as they roll off the schedule).

44) **Question 1.2.1 Technical Specifications Page 6 Capability:** How many snapshots need to be retained?

Answer: Retention requirements differ from customer to customer, due to business policy and/or legal requirements.

45) **Question 1.2.1 Technical Specifications Page 6 Capability:** What is the current file and directory count on the Isilon cluster(s)?

Answer: See the answers for questions 6 and 12 for current file and LIN counts for storage tiers A and B.

46) **Question 1.2.1 Technical Specifications Page 6 Capability:** How deep are the deepest directories/files on the Isilon cluster(s)?

Answer: Currently, we don't see any files deeper than 250 directories. However, we do not enforce any limit.

47) **Question 1.2.1 Technical Specifications Page 6 Capability:** Are snapshots being created today with "copy on write (CoW)" or "redirect on write (RoW)"?

Answer: Isilon OneFS uses both CoW (small writes) and RoW (deletes and large writes).

48) **Question 1.2.1 Technical Specifications Page 6 Capability:** How many per-file hard links are being used per cluster?

Answer: Currently, the Isilon clusters are configured for 1000 links max.

49) **Question 1.2.1 Technical Specifications Page 6 Capability:** What is the file size/count distribution for the NFS and SMB tiers, A and B?

Answer: See pg. 6, sections 4.a and 4.b for workload protocol split. As for disk capacity, system should be able to work with any ratio of SMB/NFS (99/1, 75/25, 50/50) storage. The standby site should be able to handle the same ratio should the active site fail. The remote site should not have direct client access, so access will be via backup/restore protocols.

50) **Question 1.2.1 Technical Specifications Page 6 Capability:** Do you utilize an external anti-virus and/or anti-malware scan tool?

Answer: Not currently. The desire is to support an ICAP-based antivirus scanner.

51) **Question 1.2.1 Technical Specifications Page 6 Capability:** Do you intend to use a third-party data protection tool for long term backup?

a. If so, has network bandwidth and disk performance requirements been considered?

Answer: This requirement is not part of this ITN.

52) **Question 1.2.1 Technical Specifications Page 6 Capability:** Can you define “compatible storage arrays” under item 3e on the top of page 6?

Answer: If you sell different models/products, but they share an API/protocol to allow cross-replication, then we consider them "compatible storage arrays" for the purpose of this requirement.

53) **Question 1.2.1 Technical Specifications Page 6 Capability (f):** Can you expand on section 3(f)? Is the requirement to have replication control at the granular directory folder level?

Answer: Page 6, Section 3.f is regarding failover between active/standby systems. Is this Page 8, Section 3.f? If so, this refers to the ability to use an industry-standard protocol (NFS, SMB, or NDMP) to replicate the data in that share/export off-box. This is often employed by HSM capabilities (pg. 8, 3.g) to extend the retention capabilities.

54) **Question 1.2.1 Technical Specifications Page 5 Capability (f) and (g):** Automated and/or manual failover between the 2 local sites was stated as a requirement. With an RPO and RTO of 1 hour, will a manual process using scripts be acceptable in lieu of an automated hands-off method?

Answer: The ITN requires both automated and manual failover as options. Failover must be automated, providing the option of no hands-on interaction in the event of an unexpected failure at the Active site. In all cases the RPO and RTO requirements must be met.

55) **Question 1.2.1 Technical Specifications Page 6 Capability (f):** What scenarios constitute an automated and manual failover between the active and stand-by sites as defined in item 3(f) on page 6?

Answer: Automated failover would be employed due to a power outage, network disruption, or similar unplanned event at the Active site. Manual failover would be due to planned maintenance performed at the Active site.

56) **Question 1.2.1 Technical Specifications Page 5 System Metrics (a.i):** There is a requirement for verifiable data compression and dedupe claims. What storage efficiencies are being yielded today on the Isilon for Tier A and B?

Answer: Zero percent. We do not employ data compression / deduplication on the current Isilon Tier A or B systems.

57) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** What is the file block sizes for the Tier A/B and bulk workloads for each NAS protocol?

Answer: See pp6-7, sections 4.a and 4.b, for protocol breakdowns. Protocol block sizes are assumed to be the default for each protocol (NFS=8K, SMB=16K).

58) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** What are the read/write ratios for Tier A/B and bulk workloads for each NAS protocol?

Answer: See pp6-7, sections 4.a and 4.b.

59) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** What is the connectivity and bandwidth to the existing remote site that is 300 miles away?

Answer: Current bandwidth is as much as 20Gb/sec to the existing remote site. However, as it is a shared connection, average throughput is 1-2Gb/sec.

60) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** What is the connectivity and bandwidth to the hyperscaler(s)?

Answer: If you refer to our ESXi hypervisors, then there may be as much as 64Gb/sec FibreChannel to a particular chassis.

61) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** What is the daily change rate at the primary site?

Answer: Expect Tier A to have a much higher change rate than Tier B. Estimated at 15% but as high as 40%.

62) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** How are daily, weekly and monthly backups being performed today?

Answer: There are no backup requirements associated with this ITN.

63) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** Is 54Gb/sec network requirement cumulative of all nodes in the cluster? (or per node)

Answer: Cumulative for all nodes in the cluster. Please see answers for questions 3 and 9 to understand how capacity and performance numbers were generated and how they reflect 3 years of growth.

64) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** What is the expected duration of “peak” workloads?

a. Do workloads take into consideration “backup” or “replication” activities?

Answer: Anywhere from a few seconds to a few hours during the business day. Yes.

65) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** Will production data concurrently run in the Standby Site at both locations? If not, what constitutes a failover/site switchover?

Answer: The Standby will house a hot copy of all production data. Should the Active site be unavailable, all client workloads will be redirected (switchover) to the Standby site.

66) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** When is traffic moved from Primary to Standby?

Answer: When the Active site is unavailable (unplanned emergency) or will soon be unavailable (scheduled maintenance).

67) **Question 1.2.1 Technical Specifications Page 6 System Metrics (a.i.1):** Can you list the operating System types and versions for the 20,000 connection clients in section 4.a.i.1?

Answer: No. While most UF units may be running current versions of Windows 10 or recent versions of OSX, there is no requirement for a specific OS or version for these NAS Tiers.

68) **Question 1.2.1 Technical Specifications Page 6 System Metrics (a.ii.1):** Can you list the operating System types and versions for the 7,000 connection clients in section 4.a.ii.1?

Answer: No. While most UF units may be running current versions of Windows 10 or recent versions of OSX, there is no requirement for a specific OS or version for these NAS Tiers.

69) **Question 1.2.1 Technical Specifications Page 6 System Metrics (a.i.9):** The “Tier A” median latency for NFSv4 with Kerberos in 4.a.i.9 was defined as 10ms. What is the standard deviation for the median?

Answer: New feature: no existing data available.

70) **Question 1.2.1 Technical Specifications Page 6 System Metrics (a.i.7):** For the “performance storage” read percentages defined in 4.a.i.7 on page 6, what is the ratio of sequential to random reads for the peak file system events?

Answer: This depends on the workloads.

71) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** Is there a Peak latency tolerated for NFSv4 with Kerberos or SMB3 with encryption for Tier A?

Answer: [Please refer to page 6, section 4.a.i.7.](#)

72) **Question 1.2.1 Technical Specifications Page 6 System Metrics (a.ii.9):** The “Tier B” median latency for NFSv4 with Kerberos in 4.a.ii.9 was defined as 20ms. What is the standard deviation for the median?

Answer: [New feature: no existing data available.](#)

73) **Question 1.2.1 Technical Specifications Page 6 System Metrics (a.ii.7):** For the “bulk storage” read percentages defined in 4.a.ii.7 on page 6, what is the ratio of sequential to random reads for the peak file system events?

Answer: [This depends on the workloads.](#)

74) **Question 1.2.1 Technical Specifications Page 6 System Metrics:** Is there a Peak latency tolerated for NFSv4 with Kerberos or SMB3 with encryption for Tier B?

Answer: [Please refer to page 6, section 4.a.ii.7.](#)

75) **Question 1.2.1 Technical Specifications Page 7 Capacity/Scalability:** What data set(s) make up the usable capacity for the Tier A/B and bulk workloads?

Answer: [Depends on customer requests.](#)

76) **Question 1.2.1 Technical Specifications Page 7 Capacity/Scalability:** What data set(s) will make up the 1.73TiB usable capacity at the remote site?

Answer: [Depends on customer requests.](#)

77) **Question 1.2.1 Technical Specifications Page 7 & 26Capacity/Scalability:** Does Logical Usable Capacity mean Usable Capacity under the definition 5.57 on page 26?

Answer: [Yes. We added logical to differentiate it from Physical Usable Capacity which would differ depending on how the vendors storage array is designed.](#)

78) **Question 1.2.1 Technical Specifications Page 7 Capacity/Scalability:** Capacity requirement indicate 20% growth - Is first year's growth accounted for in sizing? You are requesting a fixed price for increasing capacity assuming a 20% growth in years 2-5, totaling up to 1.5X of the original capacity. Can you confirm the assumption that we would be adding 1.26PB/Year for both the Active Site and the Standby Site? We are to assume 20% on 6.32PB per year, equaling 1.26PB per year per site? And we are to assume the same performance ratio breakdown for that capacity add? This is relative to #5 Capacity/Scalability section.

Answer: The given capacity numbers factor in current data measurements grown to the date of award, and then grown 20% per year for an additional three years. As per 5.a and 5.b, assume 6.32PB for year 1, increasing by 20%/year.

79) **Question 1.2.1 Technical Specifications Page 7 Security:** The Security requirement 6g on page 7 mentions validated FIPS140-2 standards. The Federal Information Processing Standards definition on page 23 under section 5.29 calls out level 2. Is a FIPS140-2 Security level 2 required for data at rest encryption required with item 6g on page 7? If not, what security level of FIPS140-2 is required?

Answer: FIPS-140-2, Level 2.

80) **Question 1.2.1 Technical Specifications Page 7 Security:** Is an external key management server required for any of the three locations?

Answer: No.

81) **Question 1.2.1 Technical Specifications Page 7 Security:** Does data in-flight (replication traffic) need to be encrypted between the three locations?

Answer: Yes.

82) **Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (a):** Can you define the legacy clients for NFSv3 in section 3a on page 8?

Answer: No. Any client that can only connect via NFSv3 instead of 4.0 or newer is considered a "legacy client" for NFS.

83) **Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (e):** Replication to Cloud Vendors 3(e): Data Backup, Application consuming data in the cloud, or tiering?

Answer: Data backup and/or tiering. We are not looking at cloud data consumption currently.

84) **Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (j):** Item 3(j) on page 8 calls for a client file size of 2TiB. What file block size(s) would make up this large file?

Answer: Client block transfer sizes should be assumed to be between 4K and 64K.

85) **Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (g):** Does UF have a HSM process or tool in place today? To satisfy the HSM requirement, is UF looking to tier data to a lower tier of storage based upon the age of cold data? Do you want HSM to be handle by array with object storage?

Answer: No, UF does not have an HSM process or tool in place today. Yes, UF would like to tier data based upon the access time. Yes, HSM can be handled by array with object storage, provided all other requirements can be met.

86) Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (g): What file sizes or access-times would be use cases for Hierarchical Storage Management?

Answer: HSM should work on access time of the file. We should be able to set a policy for files not accessed for a certain time period to be moved to secondary storage.

87) Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (g): Is Hierarchical Storage Management required for the active, standby or remote site?

Answer: We leave it up to the ITN respondents to determine optimal use. Note that this question was derived from the “Highly Desired” section of the ITN.

88) Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (g): With regards to the migration of data for Hierarchical Storage Management, would that include the metadata as well?

Answer: We leave it up to the ITN respondents to determine optimal use.

89) Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (g): Do snapshots need to need to be part of the Hierarchical Storage Management capability?

Answer: We leave it up to the ITN respondents to determine optimal use.

90) Question 1.2.1 Technical Specifications Page 8 Highly Desired; 3. Capability (k): What is the use case for a clone? What is your definition of clone and duration of said clone?

Answer: Potential use case: cloning a production database into a stand-alone qualitative analytics testing (QAT) tier. Depending on business processes, the clone may be required for a day, a week, a quarter, or a year.

91) Question 1.2.1 Technical Specifications Page 8 Highly Desired; 4 Security: Does the University require non-returnable drives for failures?

Answer: Not required. Data on the drive must be able to be rendered unrecoverable before disk return (degauss drive, destroy drive, destroy encryption key for SED drives, some other method). SED drives are preferred but not required if other mechanism is available that meets requirements.

92) Question 1.2.1 Technical Specifications Page 8 Highly Desired; 4 Security: Is there a WORM requirement for any sensitive data that may make up the Tier A or Tier B capacities?

Answer: No. However, the data on any returned drives must be irrecoverable. A self-encrypting drive, with the ability for UF to destroy the key before returning the drive, would meet this requirement.

93) Question 1.2.1 Technical Specifications Page 9 Highly Desired; 5 Support / Vendor:

Please clarify in item #5 Support/Vendor c. In -cloud SW updates communicated to UF Systems Administrators prior to implementation? Is this referring to firmware updates or OS updates? What specifically is this referring to?

Answer: Any version update that has the potential to impact UF's ability to access, transfer, or restore data using the supported API, must be announced in advance. This includes any potential service interruption that might occur during an update.

94) Question 1.2.1 Technical Specifications Page 9 Highly Desired; 5 Support / Vendor (d):

With regards to item 5d, what "other vendor storage systems" are in scope in addition to the existing Isilon?

Answer: None currently. However, such a re-export ability for SMB/NFS would allow us to leverage your product in assisting our customers in migrating to this new storage system from standalone legacy systems.

95) Question 1.2.1 Technical Specifications Page 9 Highly Desired; 5 Support / Vendor:

In reference to the POC required following the initial evaluation, will we be provided a test plan or could help structure the test plan of what will be evaluated? And will there be a scoring structure for that test that we can be provided visibility to?

Answer: UF will have a test plan that we will use as our guideline during any POC. It is based on activities to validate claims in proposals compared to the detailed requirements we have provided in the ITN.

96) Question 1.2.1 Technical Specifications Page 9 Highly Desired; 5 Support / Vendor

Item (c): Vendor provides support for client-performed capacity additions without requiring vendor professional services, onsite support engineer or technical support involvement. Question – Is remote/virtual support the preferred method of support, and only when necessary?

Answer: Yes. UF will determine when necessary.

97) Question 1.3 Term of Agreement Page 9:

This contract is seeking a 3year agreement up front with the ability to renew at the negotiated rate for up to (2) additional years. However, the pricing spreadsheet requests an initial cost and then recurring costs years 2-N. Do you want the initial purchase to include all file storage hardware and software and support that includes 12 months of support, then pay an annual support fee in years 2 and 3 with an option to extend that recurring support fees in years 4 and 5? Or

would you prefer to purchase 3 years of support up front, with zero recurring costs on the initial purchase, and only pay annual support costs in years 4 and 5?

Answer: Per the ITN, our intention is for an initial 3-year term with 2 optional years. However, we would like pricing specified as per Attachment A.

98) **Question 1.3 Term of Agreement Page 9:** Are leasing / finance options to be considered?

Answer: No

99) **Question 2.4 Pre-Award Negotiations Page 11:** What is the timing of this process and actual contract award and purchase? It states in 2.7 the proposals will need to be in effect for 90 days. Is that the timeline that can be expected beyond January 14th, so a final award/purchase would happen by April 14th?

Answer: There is no established time frame. It is dependent on the number of proposals received and the scheduling of any POCs.

100) **Question 4.1 Proposal Format Organization Page 16 4.1.1 Response Format Tab 9:** UF requests information on EOL equipment trade-in considerations. Is the Vendor responsible for wiping the data on this equipment or will UF perform the data wipe?

Answer: If the trade-in value is dependent upon working capacity and capability of returned hard drives, the vendor must provide the data wipe service to meet UF policies and requirements. Otherwise UF will retain the drives.

101) **Question Professional Services General:** Does UF desire implementation of storage quality of service (QoS) feature? Description: Storage QoS, can help you manage the risks that accompany meeting performance objectives for workloads. You can use Storage QoS to limit the throughput to workloads to a Storage Virtual Machine (SVM), or to groups of volumes or LUNs within an SVM, and to monitor IOPS and MBps performance.

Answer: Our review will be on the requirements of the ITN. Vendor proposals will address the methods they recommend in achieving them.

102) **Question Professional Services General:** Does UF desire implementation of Active IQ Unified Manager- Description: NetApp® Active IQ Unified Manager is a graphical management product that provides comprehensive monitoring and key management capabilities for NetApp ONTAP® systems to help manage the availability, capacity, protection, and performance risks of your storage systems.

Answer: Our review will be on the requirements of the ITN. Vendor proposals will address the methods they recommend in achieving them.

103) **Question Attachment A Pricing Proposal General:** The first column requests standard commercial pricing. Can this be list pricing?

Answer: Yes, annotate accordingly.

104) **Question Attachment A Pricing Proposal Training Cost:** How many individuals would you like to get trained? We can price it per individual, but also calculate perhaps a larger volume discount if you train a larger group of individuals.

Answer: Four Systems Administrators.

105) **Question Evaluation Criteria Page 11 Vendor Support:** "Ability to provide a cost-effective spare/replacement program. "Can you please explain what that is for? If support includes all parts replacement of the system, are you asking for the logistic process of that replacement? Or are you seeking a spare parts procurement price schedule?

Answer: This is not applicable as the part replacement shall be covered under the maintenance support agreement.

106) **Question:** When describing Protocol Op requirements please include R/W ratio and Metadata composite (what percentage of protocol ops should be considered meta) and the block size, otherwise we will assume.

Answer: Respondents will need to assume based upon their system's design. Not all systems calculate/store/use metadata the same way.

107) **Question:** We will assume that all performance requirements stated with bandwidth assume a sequential mode of operation with large block size – unless further clarification can be provided.

Answer: Block size should be assumed 8K for NFS3, 32K for NFS4, and 16K for SMB2/3. Sequential access assumption is acceptable for peak numbers.

108) **Question:** Please describe a "File System Event" (as compared to a protocol op), the context for these performance figures and how these were derived given the lack of an industry standard.

Answer: A single protocol operation (e.g., file rename) may involve dozens of file system events on the NAS system (directory open, stat, parent directory open+stat, etc.) or it may only involve one (e.g., lseek).

109) **Question:** Do the protocol ops and throughput performance requirements assume NFSv4 full encryption and SMB full encryption? If not, what percentage of the workload requires NFSv4 krb5p encryption or SMB3 encryption or signing?

Answer: Yes.

110) **Question:** Please explain the requirement; 3 (d) HTTPS-based API Object Storage

Answer: If the system offers an object storage API (e.g., HDFS), then it needs to support encryption using TLS, and UFIT needs to be able to install/manage certificates and keys.

111) **Question:** Please explain the requirement; Is this the ability to interface with an object store or do you want a proposal that includes a NAS solution and an Object Store Solution?

Answer: Page 8, section 3.d refers to a NAS solution that also offers an object-based API for client access, in addition to NFS and SMB client access.

112) **Question:** Please explain the requirement; Please expand on 4 (f) (g) in terms of the desired solution. Are you seeking anti-virus software?

Answer: Page 9, section 4.f refers to some method that your system verifies new firmware images before installing/running it. Section 4.g refers to the ability to use an ICAP server or cloud-based signature engine to ensure known-infected files are detected and/or quarantined, so the NAS cannot be used as a malware vector.

113) **Question:** What is the workflow associated with the existing system?

Answer: We have no single workflow as we have a wide variety of enterprise systems, file clients (Windows and Linux) connecting via NFS and SMB to file shares with a huge variety of configurations.

114) **Question:** What is the approximate file count for the current file storage namespace? Billions? Hundreds of millions?

Answer: See answers for questions 6 and 12.

115) **Question:** What is the approximate number of files with a size of 128K or smaller?

Answer: Tier A has over 600 M files at 128K or smaller. Tier B has over 410 M files at 128K or smaller.

116) **Question:** What are the specific services associated with one of the three cloud vendors (AWS, GCP, or Azure) that would be leveraged within UF IT?

Answer: We are provider agnostic. We are interested in using their object-based storage offerings as a replication target and/or long-term archival store.

117) **Question:** What role does file system analytics play when monitoring / troubleshooting UF file services?

Answer: Such analytics allow us to determine if reported performance issues are due to

the client network latency, authentication/authorization lookup delays, NAS system HDD IOPS issues, etc.

118) **Question:** When migrating from the existing file system, what is the cutover plan for departments, shares, and sets of files

Answer: The cutover plan will depend upon the capabilities of the new NAS system.

119) **Question:** What requirements exist to have file system events visualized with a dashboard?

Answer: It is not specified as a requirement in the ITN.

120) **Question:** Are there requirements for specific block size? How does the existing file system optimize writes to drives?

Answer: NAS protocols employ multiple block sizes. Some versions will dynamically scale according to network performance. The NAS system must be flexible enough to handle network blocks as small as 1K and up to several megabytes. The current NAS caches write data, optimizing before writing to usable storage.

121) **Question:** When drive rebuild times exceed one hour, how are the risks mitigated with the current file system?

Answer: Current cluster employs Node+Drive multiple redundancy, where each data block is reconstruct-able not only from parity data on the node, but across data from multiple nodes.

122) **Question:** What controls are in-place to avoid performance concerns when the existing file system exceeds 80% full capacity?

Answer: None. We expect performance will degrade beginning at 90%. (Existing systems exceed 80% utilization regularly; 90% is our "high-water" mark.)