



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

March 9, 2023

**ADDENDUM #1** to the University of Florida ITN23NH-121 University of Florida Information Technology Compute Environment, Support and Related Services scheduled to be opened on **March 21<sup>st</sup>, 2023 3:00 PM** 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 **ITN23NH-121** 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:**

1. Answers to supplier questions received.
2. Change in bid due date:  
Proposals are now due March 28th, 2023, at 3:00 PM ET. Submit proposals in accordance with ITN23NH-121 documents.

Sincerely,

Nicola Heredia, Director  
Procurement Services

**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

\_\_\_\_\_  
Email Address

\_\_\_\_\_  
Company Address

\_\_\_\_\_  
City/State/Zip

**Q1. 1.2.1 (4) Availability Zones: UF has identified to data centers for the solution, does the 2 availability zones refer to the two data centers or does each data center require 2 availability zones for a total of 4?**

A1. UF availability zones span datacenters. Half of the hosts in each Availability Zone come from each datacenter, so the loss of a datacenter results in only a 50% capacity loss.

**Q2. 1.2.1 (4) Availability: ITN request 50-80 Nodes in each availability zone does this refer to each data center?**

A2. Availability Zones span datacenters, so the 50-80 number figure should be x 2 to cover both Availability Zones.

**Q3. 1.2.1 (4) Availability: ITN request 50-80 Nodes what is the desired qty of nodes?**

A3. UF does not have a specific quantity in mind – it's up to the vendor to propose a solution that meets the compute specifications specified in the ITN.

**Q4. Can UF get clarify to what exactly are the authentication requirements against the Duo Cloud with the winning on-premises solution?**

A4. The solution should be able to use Duo Cloud as a second factor for authenticating the users

**Q5. Hardware technology has progressed to a point that power and heat from next generation processors, memory, solid state drives, and other server components are causing server manufacturers to redesign their products. The blade technology focuses on integrating these components in a compact platform that present physical challenges in the datacenter. Lenovo has made decision to move away from the blade technology and focus on overcoming these challenges on our rack, edge, liquid-cooled, and mission critical servers. Would the RFP be open to presenting an alternative solution to blade technology?**

A5. Yes – UF is open to alternate solutions as well as blade-based solutions.

**Q6. Are we able to run a Live Optics capture on the existing environment in scope?**

A6. No

**Q7. Are you able to provide the most recent RVTools export?**

A7. Please see Attachment C for deidentified data.

**Q8. Is the environment in scope all virtual or including physical servers as well? If mixed, are you able to provide details / clarity?**

A8. The environment mostly (>95%) consists of compute used for virtual servers. There are a small number of servers that are bare metal which are used for Oracle or management infrastructure. This ITN is for refreshing the virtual servers.

**Q9. Are you able to share any Visio drawings – high level?**

A9. <https://hosting.it.ufl.edu/services/virtual-machine-hosting/>

**Q10. How many different clusters are part of the environment in this scope of work today?**

A10. 5 clusters are in scope for this ITN.

**Q11. From the UF requirement of a blade solution, it looks like most blades will look to be booted from the local SSD RAID-1 configured SSD drives. Thus, will all workloads leverage external storage or will we need to architect a vSAN solution as well?**

A11. The workloads will leverage fiber channel or NFS based external storage. While UF currently uses NFS and FC, we are open to vSAN based solutions.

**Q12. What databases / workloads are being utilized today? SQL, Oracle, VDI, HPC, etc.?**

A12. UF IT serves as a hosting provider for departments and colleges on campus. Most VMs are general mixed use server hosting, with some Oracle, MySQL, and MS-SQL hosting.

**Q13. If VDI is in play here, are we able to get the details on the number of users, type, applications, existing details on performance, user storage needs, connectivity, Horizon/Citrix, GPU requirements, etc.?**

A13. VDI is not part of this ITN

**Q14. Does UF plan on reusing existing VMware licensing or will we need to include it as part of the new solutions? (ex. vSphere Enterprise Plus, vCenter Standard, Horizon licensing, Aria Automation (vRA) / Operations vROps, etc.)**

A14. UF currently has a ELA with VMWare that will be utilized with this environment but are open to other licensing agreements if it makes financial sense.

**Q15. Does UF currently utilize NSX within the environment? If so, any logical network, extended L2 between the existing two data centers today?**

A15. NSX is not currently in use.

**Q16. If new VMware licenses are required to be included, is it also preferred/required to have with 5 years of support?**

A16. Yes

**Q17. What is the current storage that is connected to the HPE Synergy today? Is it mostly FC?**

A17. 90% fiberchannel, 10% NFS.

**Q18. What FC SAN fabric / switches are leveraged today?**

A18. Cisco MDS 9000 series

**Q19. Are the existing availability zones (Data Centers) collocations or UF on-premise data centers?**

A19. The Availability Zones are spanned between an on-campus datacenter and another UF owned datacenter in the metro area that are connected by redundant private fiber.

**Q20. Is there additional, existing power / cooling, as well as additional rack space to accommodate the new solution, while migrating from the old one, at each of the data centers?**

A20. There is sufficient power and cooling reserve and additional rack space available to host the new solution in order to migrate.

**Q21. Is the requirement for an Active / Active design or Active / Passive?**

A21. UFs current design is Active/Active.

**Q22. Does UF leverage a Spine / Leaf network today or traditional architecture?**

A22. The network is spine/leaf in design.

**Q23. What is the existing network today? Are you able to share Visio diagrams?**

A23. The current network is comprised of multiple 10GB links aggregated to a total of 80GB/rack. There are 4 racks in the current implementation. 2 in each data center.

**Q24. Is there a preference of UF on whom does the professional services, OEM or a VAR?**

A24. Professional services are not required for this deployment.

**Q25. Will UF need professional services with workload migration as well? If so, please include details of all workloads, applications, ESXi versions, etc.)**

A25. Professional services are not required for this migration.

**Q26. If NSX is in play, will UF need professional services to deploy / configure that in a new environment? If so, please provide details on the requirements.**

A26. UF does not utilize NSX.

**Q27. Will UF require managed services for the new solution?**

A27. UF will not require managed services for the new solution.

**Q28. The ITN document did not list, but what is UF leveraging for backup today to protect the workloads? (S/W, H/W, tiering to cloud, etc.,)**

A28. This is out of scope for the ITN.

**Q29. Will you need professional services to configure & integrate the new environment or will UF do that? If so, we will need details around workloads, number, backup size, rate of change, retention, tiering, etc.**

A29. UF does not need professional services to configure and integrate the new environment.

**Q30. Does the solution require that all cables be included? (ex. OM4 LC-LC, AOC / break-out, DAC Twinax, etc.) If so, are there any length requirements?**

A30. UF will provide all cabling required to connect both storage and ethernet networks to the rest of the data center. Vendor is to provide any cabling required within the solution such as frame or chassis stacking cables. Vendor should also include quotes for required server-side optics.

**Q31. What is the minimum amount of storage per blade required?**

A31. 240G

**Q32. What are the performance requirements of the workload?**

A32. Local storage will only be used for ESXi OS deployment or in some cases for bare metal deployments

**Q33. Have you run an assessment on your environment?**

A33. RVtools output will be provided.

**Q34. The ITN document states 50 to 80 nodes per availability zones. How many nodes are you looking for in the initial configuration / design?**

A34. UF is flexible on the number of nodes included provided that the compute specifications are met.

**Q35. If the solution requires up to 80 nodes per availability zone, would you need a fixed amount of chassis (# of blade slots) to support existing initial configuration or all of the chassis to support the growth?**

A35. Initial slot capacity should accommodate initial purchase plus an additional 20% growth.

**Q36. (ITN Document page 7, "5. Graphics - a. GPU support connected direct to blade.") Does UF have any requirements for a specific card model?**

A36. Not at this time

**Q37. It's not mentioned in the ITN document: is there any backup software or strategy that needs to be included?**

A37. Backups are not in scope for this ITN

**Q38. Do you want the Excel Attachments A & B to remain in Excel for the submission?**

A38. Yes

**Q39. In attachment B, Compute Feature List.  
10. Storage**

**a. Capability to have 2-8 NVMe or SSD Drives Available to each blade.**

**We do not see a SSD size requirement, is there a sizing requirement or just the capability to support up to 8 NVMe or SSD drives per blade?**

A39. The compute hardware would ideally support at least two NVMe or SATA/SAS SSD drives to be able to boot ESXi or other OS. If vSAN is proposed, then the additional drives and capacities will need to be determined as part of the solution.

**Q40. Is a blade solution required? Architecture states it's a required solution but under Networking (8.e) states if bladed solution?**

A40. No

**Q41. Is ESXI the only hypervisor utilized for the entire solution set?**

A41. Yes, at this time.

**Q42. Any additional physical hosts required to be rolled into Orchestration layer?**

A42. UF is unsure of the meaning of "Orchestration layer" for the proposed solution.

**Q43. How did you determine 50-80 nodes for each Zone? Is this flexible?**

A43. The current design consists of two Availability Zones with 50% capacity from each datacenter in each Zone. The 50-80 number is based on current infrastructure and is meant to include N+1 compute per cluster in each Availability Zone at under 50% utilization for CPU and Memory use. UF is flexible on this number if the design can meet the compute specifications. This number does not accommodate for 20% future growth.

**Q44. What is the average latency between the two AZ's?**

A44. <2ms round trip time.

**Q45. Is there a requirement for Stretched Layer-2 or Traditional Layer 3 between AZ's?**

A45. Layer 3 separation is required between Availability zones. As a reminder an AZ can span multiple data centers. Layer-2 within an AZ is a requirement.

**Q46. Is it the intent to provide compute functions based on workload domains?**

A46. No, as a hosting provider for campus departments and colleges, there are mostly mixed workloads.

**Q47. Can you share any workload domain layout? (i.e., General Compute, HPC Clusters, EUC, Containers & AI/ML)?**

A47. Mostly General Compute though there are small amounts of EUC. Containers are in use but deployed via VMs. There is interest in Kubernetes.

**Q48. Will the virtualization platform be traditional VMware or VMware VCF (VMware Cloud Foundations)?**

A48. Traditional VMWare but we are open to alternatives.

**Q49. Is N+1 required for patching and firmware updates? No listing of "Acceptable Downtime"**

A49. Yes. VMWare clusters are designed as N+1. Availability Zones are designed with 50% redundancy to sustain a whole datacenter failure.

**Q50. What is the current networking hardware make/model/speeds at the network aggregation layer?**

A50. Current connections are running aggregates of 10GB connections, 8 per rack for 4 racks. They also support 25 and 100gb connectivity, which is planned to be used for the new solution.

**Q51. What is the current storage networking hardware make/model/speeds at the storage network aggregation layer?**

A51. The current main is utilizing Cisco fiber channel switches. 8 – 8GB fiber channel connections go to each rack. 16/32G FC is preferred for this solution.

**Q52. Is there a current GPU Manufacturer Make/Model determined for GPU support? If so please share make/model for proper power/cooling considerations.**

A52. Not at this time.

**Q53. Is there a preference for Outright Purchase or Consumption as a Service?**

A53. No.

**Q54. Will this solution require ROCE networking support? (RDMA over Converged Ethernet)**

A54. No.

**Q55. Will the local storage requested be used for VSAN? If not, what is their primary purpose?**

A55. Local storage will be utilized to boot the OS. If vSAN is part of the proposed solution then additional disks will be required.

**Q56. Will future workloads be deployed via Terraform, Ansible or VMware Aria Automation?**

A56. All three of these deployment alternatives are currently being utilized in one manner or another to deploy VMs on the infrastructure.

**Q57. Due to various constraints, should the solution be solely based on Intel CPUs or are there other acceptable options?**

A57. Intel or AMD processors would be considered for this solution. Homogeneous use of a CPU vendor and CPU family is a requirement for initial deployment/purchase.

**Q58. What are the current power configurations within the Datacenter (20am or 30amp) (any limitations on HighVoltage vs LowVoltage 120/240)?**

A58. The datacenter utilizes high power bus bar infrastructure to deliver power to the racks.

**Q59. Do you have a planned award date and/or a deployment timeframe for this project?**

A59. Award data and deployment timeframe are to be determined. Lead times should be included with all proposals.