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

November 5th, 2019

ADDENDUM #1 to the University of Florida ITN20RL-121 CORE Network Equipment Refresh scheduled to open on **Thursday November 14th, 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 **ITN20RL-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:**

•Answers to questions asked prior to the deadline of 5pm, October 29th, 2019

Sincerely,

Rob Luetjen

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

Vendor Questions ITN20RL-121 CORE Network Equipment Refresh

1) **Question: Are you doing MPLS to Edge or TOR currently, or is it a future requirement?**
   Answer: We are doing MPLS to the edge now. We are also employing 4 route reflectors. We are not using any "TOR" devices. They are independent layer 3 switches that feed each building.

2) Question: Are you doing EVPN to Edge or TOR currently, or is it a future requirement?

Answer: No current EVPN. This is a future possibility.

3) **Question: What is timeframe for 400G requirement?**
   Answer: That depends on a few items:
   - Support for "FlexE" (Flexible Ethernet) may delay our needs for 400G
   - Support for 200G using QSFP56-LR optics may delay our need for 400G
   - Based on current growth rates and assuming no support for the above items, I would estimate our need for 400G in the 2022 timeframe.

4) **Question: Is VPLS a Day-One requirement or is EVPN-MPLS suitable?**
   Answer: Currently our edge platform doesn't support EVPN-MPLS directly, however our desire is to move toward EVPN. In the interim, extending vlans from the building to the core and using EVPN-MPLS will be suitable.

5) **Question: What model of Nexus switches are currently being used at both campuses?**
   Answer: These are not "Cisco Nexus" switches. We are using Cisco 6509 Sup2T in the "Nexus" role. Don't confuse the "Nexus" role with the Nexus brand of switch. Our building networks currently use Cisco 3850 and Cisco 9500 switches which will be connecting to the new core network

6) **Question: Can you provide a detailed diagram listing the connections between the TOR switches and the Catalyst 6506/6509?**
   Answer: Each core router has 30-50 connections to buildings, at rates between 1G and 10Gbps. The edge layer 3 switch is currently a 3850 or 9500 (some limited 3750s, but those are going away). These switches are either straight layer 3 / OSPF or MPLS back to the core. There are no "TOR" switches per se. Each building has 2 connections back to two independent core routers (layer 3 or MPLS). Please see the example building level diagram below:

Wallplate Network

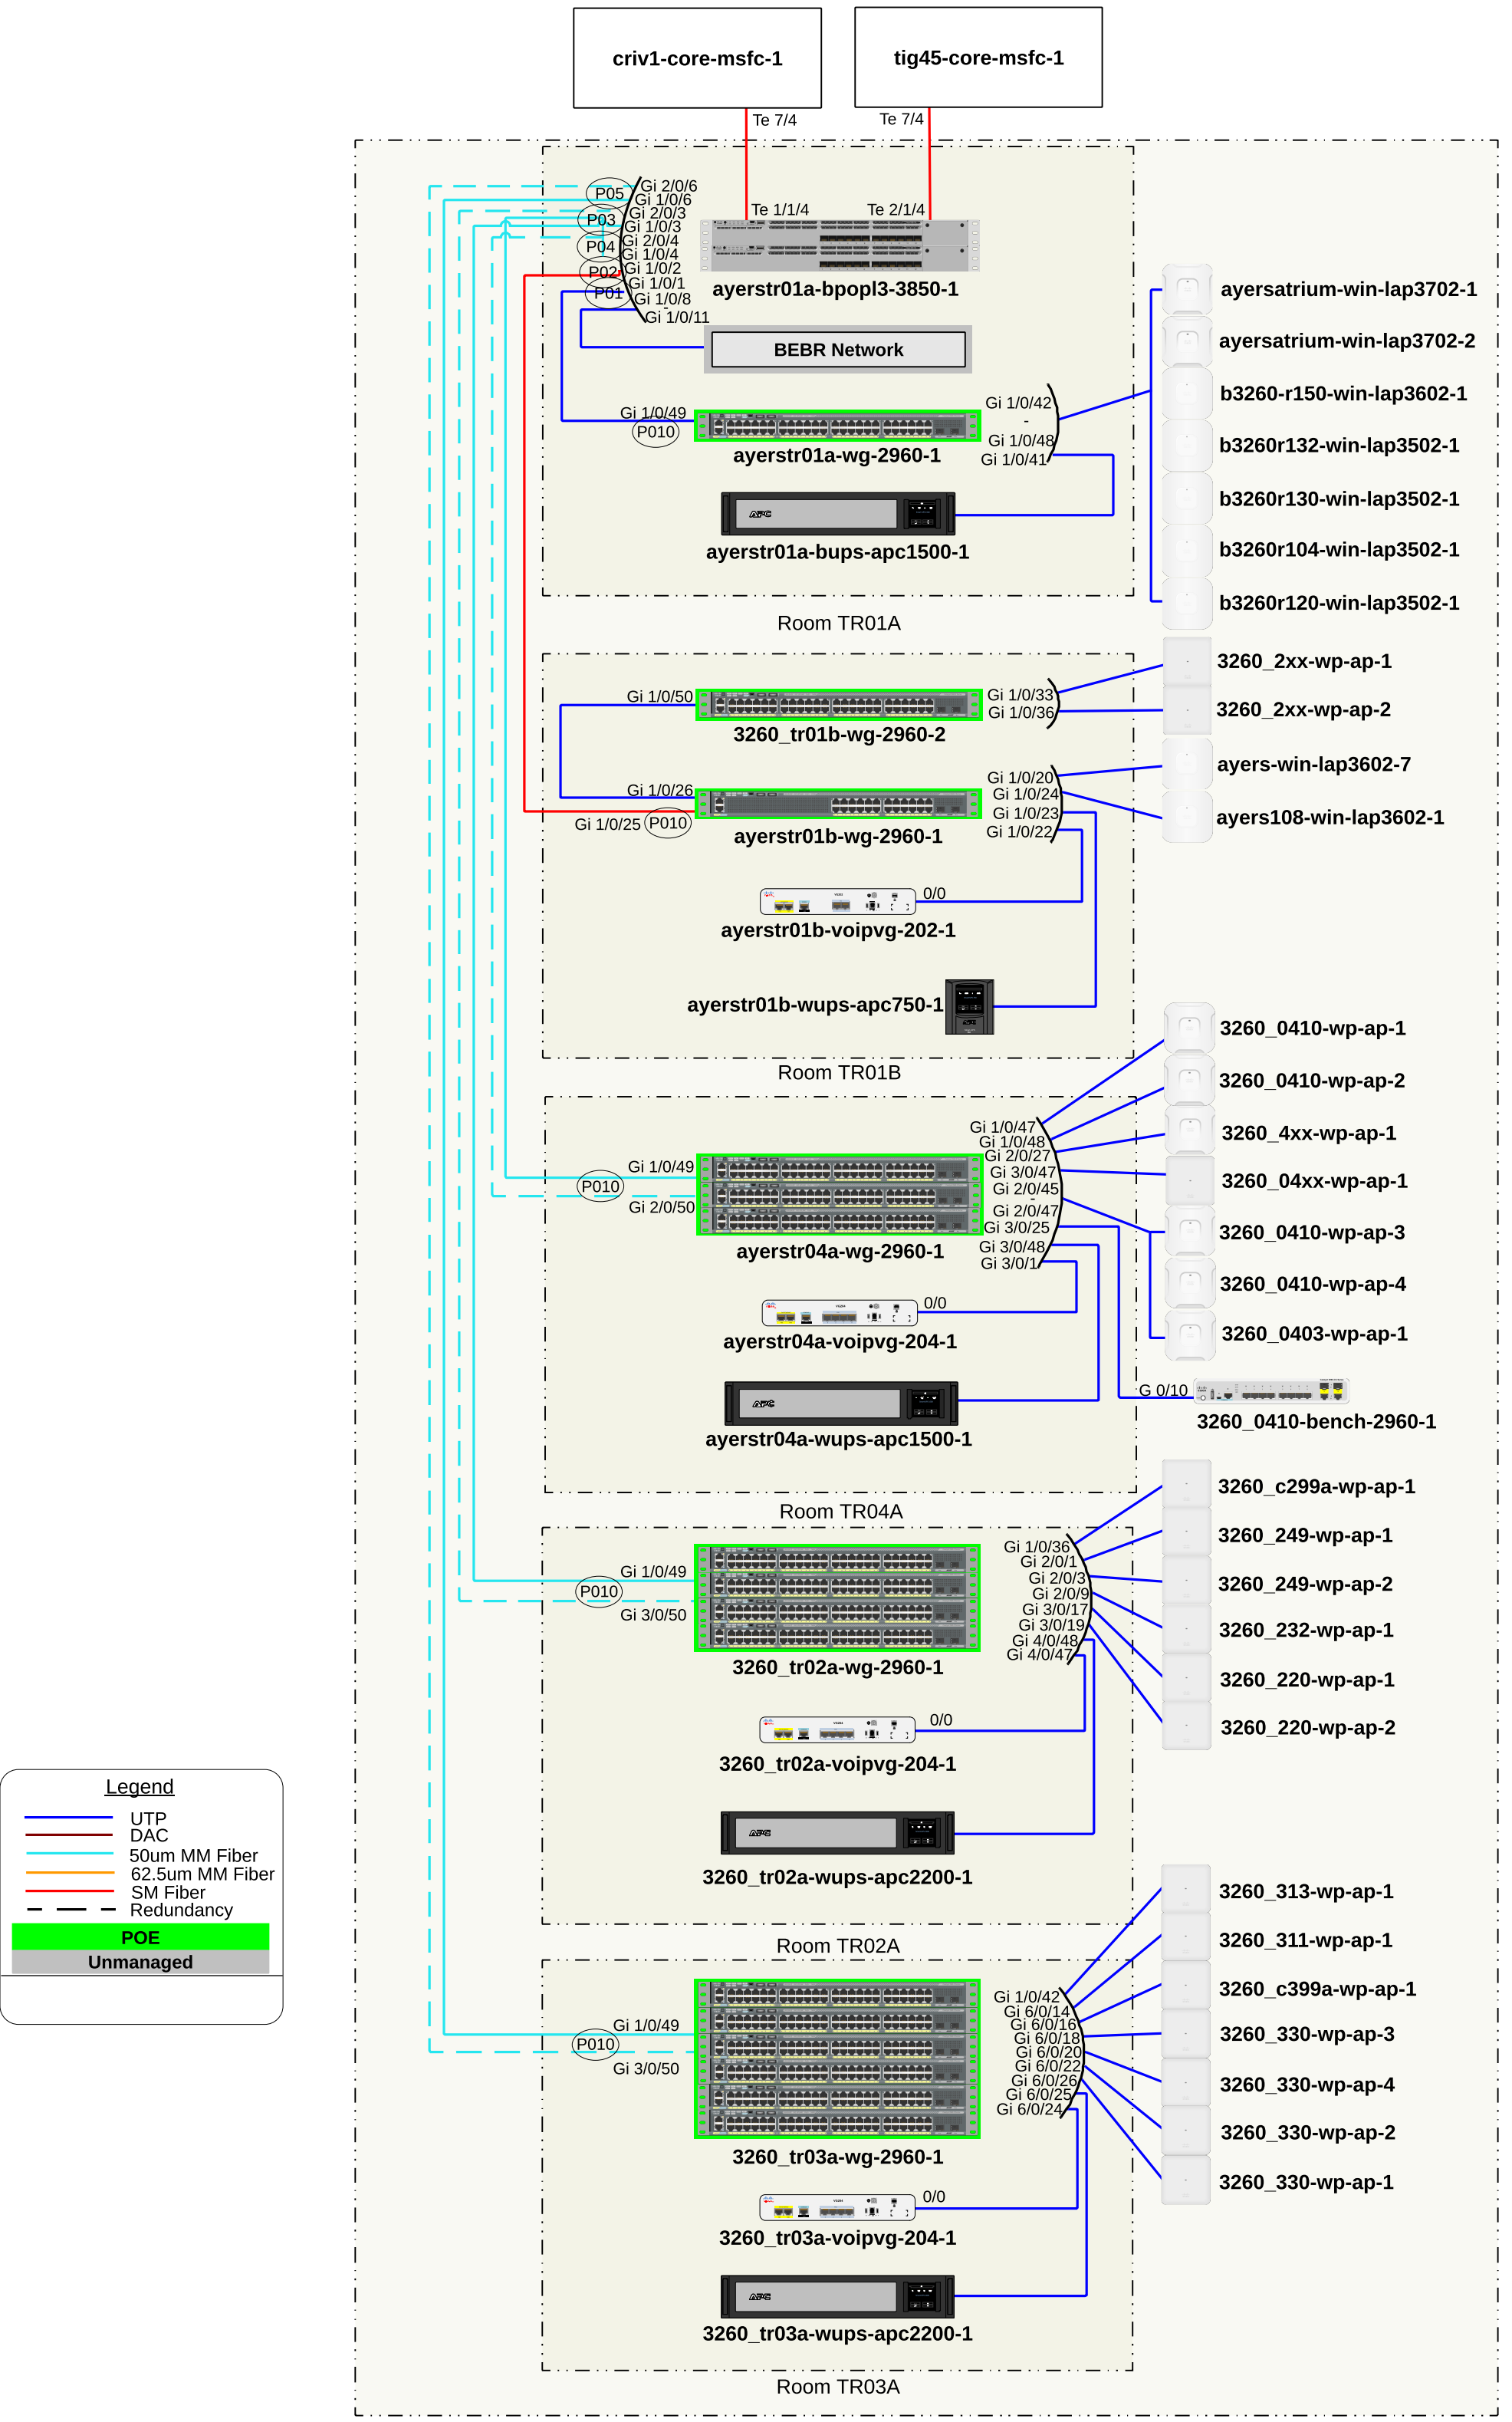

Legend

- UTP
- DAC
- 50um MM Fiber
- 62.5um MM Fiber
- SM Fiber
- Redundancy

POE

Unmanaged

7) **Question: Can you provide a detailed diagram listing the connections between each Nexus?**
   Answer: The complete core router network diagram (interconnections between core routers were provided as part of the ITN packet). Again there are no “Cisco Nexus” devices in use. Our core network has an “outer core” which is 14 Sup2T 6509s that each branch out to the buildings and 4 “inner core” devices which connect all the outer core routers together. These are also Sup2T 6509 routers.

8) **Question: Can you provide a detailed diagram listing the connections coming from the Nexus going upstream?**
   Answer: The “Core-Nexus” layer connects to our Juniper MX10008 Ewan routers.

9) **Question: How are Nexus switches interconnected physically and logically?**
   Answer: The core routers route between each other using a 802.1q trunked Ethernet connection. These are also MPLS enabled. Nearly all of our backbone traffic is MPLS encapsulated.

10) **Question: Is there a need to support both EVPN-MPLS and EVPN-VXLAN within same hardware?**
    Answer: At this point, it is unlikely EVPN-VXLAN termination will be required at the core layer. If given the option, we would prefer EVPN-MPLS.

11) **Question: Can you provide us a detailed port count on the Catalyst 6506, 6509s, and Nexus switches?**
    Answer: Note that this port count does change fairly frequency (goes up) as new buildings are built or upgraded so some additional margin for demand growth + margin for future growth after implementation is recommended.

    ssrb230a core 6509: 23 1G, 24 10G
    csev1 core 6509: 28 1G, 19 10G
    mebv1 core 6509: 22 1G, 21 10G
    ctx36 core 6509: 17 1G, 23 10G
    criv1 core 6509: 20 1G, 24 10G
    tig45 core 6509: 17 1G, 21 10G
    elmv1 core 6509: 22 1G, 18 10G
    aerv1 core 6509: 19 1G, 18 10G
    mowv1 core 6509: 15 1G, 10 10G
    fifv1 core 6509: 25 1G, 9 10G
    smpb151 core 6807: 22 10G
    mbig161 core 6807: 22 10G
    ecdctr01a core 6506: 7 1G, 12 10G
    escv1 core 6506: 7 1G, 12 10G
    ssrb230a nexus 6509: 36 10G
    ctx36 nexus 6509: 1 1G, 36 10G
    escv1 nexus 6506: 12 10G
    ecdctr01a nexus 6506: 12 10G

12) **Question: Do the port requirements listed in the RFP apply to the core ring routers and 6500 switches or just 6500's?**
Answer: If an inner core ring is maintained, the ports will primarily be 100G, but the need for 48x100G would be reasonable.

13) **Question: The RFP says ability "Ability to grow to at least 48 x 100G via QSFP28 / 192 x 10G per chassis. 10x10 breakout solutions are preferred, however 4x10 or "satellite" switch configurations will also be accepted [2]". The RFP doesn't specify the exact port counts needed for the breakouts. What is the starting port count required? Is 48 x 100G / 192 x 10G an inclusive requirement?**
Answer: The exact number of ports per chassis will depend on how many chassis you are proposing. The current port count for the current network is listed below. Note that this port count does change fairly frequency (goes up) as new buildings are built or upgraded so some additional margin for demand growth + margin for future growth after implementation is recommended.

ssrb230a core 6509: 23 1G, 24 10G
csev1 core 6509: 28 1G, 19 10G
mebv1 core 6509: 22 1G, 21 10G
ctx36 core 6509: 17 1G, 23 10G
criv1 core 6509: 20 1G, 24 10G
tig45 core 6509: 17 1G, 21 10G
elmv1 core 6509: 22 1G, 18 10G
aerv1 core 6509: 19 1G, 18 10G
mowv1 core 6509: 15 1G, 10 10G
fifv1 core 6509: 25 1G, 9 10G
smpb151 core 6807: 22 10G
mbig161 core 6807: 22 10G
ecdctr01a core 6506: 7 1G, 12 10G
escv1 core 6506: 7 1G, 12 10G
ssrb230a nexus 6509: 36 10G
ctx36 nexus 6509: 1 1G, 36 10G
escv1 nexus 6506: 12 10G
ecdctr01a nexus 6506: 12 10G

14) **Question: What are the port requirements for replacing the Nexus core ring routers? Is the Core Nexus just a name and doesn't imply that they are Cisco Nexus switches?**
Answer: Yes, there are no Cisco Nexus switches in use. The name "nexus" predates the Cisco Nexus line. The exact nexus port counts for the current network are listed above as "nexus 6509" or "nexus 6506" devices. Note that in the new design, all core to core links and core to nexus links will be 100G.

15) **Question: Is redundant RP a must have for core-ring and core?**
Answer: Either redundant RP, or two boxes in the same location with full state exchange (virtual chassis)

16) **Question: Will the new core have 100G links to the core ring?**
Answer: Yes

17) **Question: Does the 10x10 or 4x10 breakouts need to be SR or LR?**
Answer: Mostly LR. We would like to be able to support some SR for local connections but can use all LR if needed.

18) **Question: What type of fiber and distance exists between the core-ring and the core devices?**
Answer: All connections are within spec for a 10G LR or 100G LR4 optic. All connections are currently running on 10G LR.

19) **Question: With the satellite solution, what kind of distance is expected between the satellite and the far end device?**
Answer: No more than 1G LX, or 10G LR.

20) **Question: Is it expected for satellite devices to support 1G and 10G?**
Answer: It would be preferable for the satellite device to support 10G as well. It will be required if the primary chassis is proposed as a 100G only solution. It is also highly preferred that any satellite device be “virtually” integrated into the larger chassis to reduce the number of individual elements we have to maintain.

21) **Question: Is the ctx36-nexus-msfc-1 and ssrb230a-nexus-msfc-1 in the main campus or east?**
Answer: The only devices on east campus are:

ecdctr01a **
escv1 **

Two core devices and two core-nexus devices.

22) **Question: What kind of connection exists between the 6500 and the core ring routers (LAG, MC-LAG, L3, MPLS, L2)?**
Answer: The are MPLS based routed vlans. There is currently no LAG or MC-LAG. The interface is configured as a trunk to allow the passing of vlans if necessary, but all traffic is MPLS Layer 3 routed.

23) **Question: Does the following statement imply that the control plane synchronization is needed for both data + control between chassis? "Redundant Route Processor and/or multi chassis stateful configurations [3]"**
Answer: If one box per location is provided, only redundant RPs are necessary. If two boxes per location (i.e. virtual chassis), then both boxes must be fully synchronized. Think something along the lines of Cisco VSS. We will not be doing VSS across locations but are willing to look at multi-boxes in the same location to provide RP level redundancy.

24) **Question: Section 4.1.3 appears to be missing in the ITN document. Please confirm that this is just a paragraph numbering issue or if something is missing and needs to be provided.**
Answer: This is a numbering error.

25) **Question: Understanding that one of the requirements is for a Fully distributed forwarding plane with distributed on-demand ACL support, if all other requirements can be met, is a solution with a centralized forwarding architecture acceptable?**
Answer: We would need to fully analyze the scale of the proposed solution to know if it would be acceptable. You may provide this as a possible solution, but if the scalability is in doubt, it may be discarded.

26) **Question: Will the new Core Network devices be deployed as MPLS P routers? Or PE as well, at least for transition?**
Answer: All devices should be able to run in both P and PE mode, just for flexibility. Architecturally, we would like all core routers to be P routers, but there are several instances where this might not be possible soon or at all.

27) **Question: In regard to the overall integration, what are the existing E-WAN / Internet Edge routers? Is there a current plan to migrate to something else, and if so, what platform?**
Answer: The existing E-Wan routers are currently connected to the ssrb230a-nexus-msfc-1 and ctx36-nexus-msfc-1. During the refresh, the type and number of connections will be dependent on the topology proposed. We will be connecting the existing Ewan routers to at least two other core or core-nexus routers. The EWan is a new platform and not planned to be replaced.

28) **Question: Would UF accept a CFP-based 100G ports/line cards?**
Answer: No

29) **Question: Envisioned use cases for Container/VM/Plug in module support?**
Answer: Distributed packet analysis (wireshark). Distributed testing (perfsonar), etc.

30) **Question: There is reference to having the ability to grow to 48 ports of 100G / 192 ports of 10G, and also reference to support for 1G interfaces. What per router port density should be built into the solution? What is the minimum port count for 100G, 10G and 1G?**
Answer: It is envisioned that 1G support would be handled via a sub-tending shelf to be able to provide as high a density solution as possible in the primary shelf.

The current in-use port count is as follows:

ssrb230a core 6509: 23 1G, 24 10G
csev1 core 6509: 28 1G, 19 10G
mebv1 core 6509: 22 1G, 21 10G

ctx36 core 6509: 17 1G, 23 10G
criv1 core 6509: 20 1G, 24 10G
tig45 core 6509: 17 1G, 21 10G
elmv1 core 6509: 22 1G, 18 10G
aerv1 core 6509: 19 1G, 18 10G
mowv1 core 6509: 15 1G, 10 10G
fifv1 core 6509: 25 1G, 9 10G
smpb151 core 6807: 22 10G
mbig161 core 6807: 22 10G
ecdctr01a core 6506: 7 1G, 12 10G
escv1 core 6506: 7 1G, 12 10G
ssrb230a nexus 6509: 36 10G
ctx36 nexus 6509: 1 1G, 36 10G
escv1 nexus 6506: 12 10G
ecdctr01a nexus 6506: 12 10G

31) **Question: Of the 48 ports of 100G / 192 ports of 10G, what percentage of ports do are expected to be 10G vs 100G?**
Answer: We prefer that all ports on the primary chassis be 100G and usable as either 4x10 or 10x10 with a breakout. We expect the 100G port count to be very low to begin with, but we would like to be able to migrate without adding much additional hardware over time. Another option is 100G only on the primary box and a "virtual linecard"/satellite which supports 1G and 10G. In this mode, all of the 1G ports should be multi-rate 1G/10G.

32) **Question: Should optics/transceivers be included in this response? If so, for all ports?**
Answer: Please include them as an option. You may assume from the port counts above that 1G is 1G LX and 10G is 10G LR.

33) **Question: Would a mixed box solution (from same vendor/manufacturer) be acceptable?**
Answer: Yes. It is preferred that this mixed box solution be integrated as a single virtual chassis, however.

34) **Question: Relating to these requirements:**
    - **Support for policy routing at scale (at least 64K entries)**
    - **Support for static routing at scale (at least 64k entries)**
    - **Support for traffic ACLs at scale (at least 128K entries)**

Could you clarify the use case(s) for the scale requirement of ACLs, PBR, static routes?
Answer: Policy and static routing use cases are for cases where we are implementing above the network security and need to redirect traffic for analysis. This may involve also dropping traffic as close to the host as possible. If scale is an issue, please provide your scale numbers and we will determine if its still workable.

35) **Question: Can you clarify what is meant by "full width" in the requirement for IPv6 support with “full width” IPv6 ACL support?**
Answer: No “address compaction” when implementing ACLs. The ACL should be able to express the full 128-bit address and not some combination of the first part of the address and the last (i.e. first 48 bits and last 16 bits for example).

36) **Question: For any requirements marked with [3], if no testing has been done to validate "the degree that the given specification exceeds the minimum," will we be able to work with UF on testing to validate the degree? And what is the recommended way to respond to the ITN when there is no test data?**
Answer: All items marked with a 3 will be tested, either in vendor labs or in UF labs depending on equipment proposed. We will work with the vendor to make sure all testing is 100% accurate and provides a true understanding of the performance of the equipment.

37) **Question: Could you clarify what is meant by the requirement Service Provider oriented feature set?**
Answer: UF operates its core network very much like a service provider would. We employ common service provider feature sets such as MPLS, BGP, VRFs, MVPN, high internal scale, etc. We prefer the equipment suggested are developed with service providers in mind. We understand that often Enterprise and SP feature sets do converge, we just want to make sure that the equipment that is suggested is done so with SP needs in mind.

38) **Question: Is there a way to submit sensitive performance information in a way that is not shared publicly beyond this ITN?**
Answer: Per section 4.2.18 of the ITN document, if the vendor needs to submit proprietary information with the proposal, the vendor shall ensure that it is enclosed in a separate envelope from the proposal and that it is clearly designated and conspicuously labeled as such. Vendors who submit responses with information noted as proprietary may be asked to substantiate why the information is proprietary or is otherwise exempt from a public records request under Florida Law.

39) **Question: GP PIC support, IPv4 only or IPv6 as well? And BGP PIC Core or PIC Edge or both?**
Answer: V4 with V6 on roadmap. Both BGP PIC core and edge.

40) **Question: What are the current VPLS use cases?**
Answer: VPLS use cases are to provide internal layer 2 only network to some networks on campus. These may include distributed research devices which must remain private, networks where compliance is an issue (i.e. FISMA, CJIS, etc). Other technologies to provide multipoint services such as EVPN are also acceptable. We would only use VPLS if it included BGP based signaling (RFC 4761)

41) **Question: Is the highly desired EVPN-VXLAN support requested as a potential VPLS replacement? If not, what is desired used case for VXLAN EVPN?**

Answer: Yes. Additionally, EVPN-MPLS is also another possible replacement for VPLS.

42) **Question: For the requirement for Support for traffic ACLs at scale (at least 128K entries), is this for 128K different ACLs or 128K ACEs?**
Answer: ACEs

43) **Question: For the requirements for Support for policy routing at scale (at least 64K entries), Support for static routing at scale (at least 64k entries) and Support for traffic ACLs at scale (at least 128K entries), the requirements for this scale to be system-wide? Or per line card? Or something else?**
Answer: System wide, however the scale demand may be focused within the chassis, so a distributed solution should be able to support this or similar scale on a per line card basis.

44) **Question: For the requirements for Support for 1G interfaces, would this fiber 1G only? Or copper 1G as well?**
Answer: Almost all fiber, but selective copper is a possibility (adjacent devices).

45) **Question: For the requirement for Future support / Growth path to 400G interfaces, what quantity of 400G ports per chassis is anticipated?**
Answer: We expect 400G to be purely for uplink. The exact port count depends on the solution provided, but we do not expect to extend 400G into the buildings. The exception to this would be our HPC center at East Campus. If a “core nexus” ring is proposed, additional 400G interfaces should be planned to feed the HPC center.

46) **Question: For the requirement for Fast RIB to FIB/TCAM programming, is the questions specific to RIB to FIB programming? It is unclear how TCAM relates.**
Answer: This is specific where TCAM is used in the forwarding path, which is a common configuration. We will be testing to determine how quickly the hardware forwarding path will be programmed from the RIB/control plane.

47) **Question: What is the anticipated purchase and deployment timeframe for this solution?**
Answer: The exact answer is dependent on the solution proposed. Two possible options:

1. All equipment is procured in FY19/20 and deployed during FY19/20 with some work occurring early in FY20/21 (CY20).

2. Half of the equipment is procured and deployed in FY19/20 and the other half in FY20/21.

48) **Question: What is the physical connectivity from main campus to east campus? Is UF open to pulling new fiber to support more redundancy to east campus?**

Answer: UF has redundant dark fiber between main and east campus. One side of the ring is heavily used and its possible we may deploy additional fiber. This would be determined after analyzing the proposed solution.

49) **Question: What is considered fast in reference to fast TCAM programming etc? Is UF looking for measured times?**
Answer: Yes, all items marked with a 3 will be tested either in UF labs or in vendor labs. We will be measuring the time it takes for prefixes to be programmed into hardware.

50) **Question: Can you define "full width" in reference to IPv6 ACLs?**
Answer: No "address compaction" when implementing ACLs is acceptable. The ACL should be able to express the full 128-bit address and not some combination of the first part of the address and the last (i.e. first 48 bits and last 16 bits for example).

51) **Question: Support for 1G interfaces, is this 1G-BaseT and/or 1G SFP interface?**
Answer: Most 1G is fiber, but we would like to be able to do 1G over copper occasionally.

52) **Question: Will UF accept sFlow in addition to Netflow/IPFIX?**
Answer: sFlow is less preferred for our core use case. Netflow/IPFIX is highly preferred.

53) **Question: Please provide context around policy route scale requirements?**
Answer: Policy and static routing use cases are for cases where we are implementing above the network security and need to redirect traffic for analysis. This may involve also dropping traffic as close to the host as possible. If scale is an issue, please provide your scale numbers and we will determine if its still workable.

54) **Question: Please provide context around ACL scale requirements?**
Answer: We implement ACLs on the core as a supplement to our bpops where we do most of our layer 3 routing. Scale is occasionally needed for sudden security issues. Additionally, we are working on security automation which may need to push campus rulesets out to the network at scale.

55) **Question: Is the VOIP multicast used exclusively for emergency notification?**
Answer: VOIP multicast is, but not all multicast. Multicast is also heavily used for cloning/ghosting or machines, research use, etc.

56) **Question: Please provide context on why mVPN is currently implemented? What problem is mVPN solving for the University?**
Answer: Our backbone is MPLS and VRF based. To provide multicast over such an infrastructure and maintaining the VRF separation, we must use mVPN.

57) **Question: What determines the membership of devices in the multicast implementation?**

Answer: Nearly all networks at UF are multicast enabled. Most multicast use is currently ASM. There is some SSM, but vendor support is very limited. Membership is specific to multicast joins sent from the hosts.

58) **Question: How many VRFs are required?**
Answer: We would like to support 256 VRFs max. If this number cannot be met, please provide your scalability and we will determine if this is an issue or not.

59) **Question: Where does the network segmentation originate? I.e. Where are the VPNs/VRFs originated?**
Answer: VRFs origination is a combination of our core being PE devices, or PE devices within the buildings. The newer buildings are moving to MPLS inside the building, and this is our overall direction, however it will take a while to get to that point.

60) **Question: Where does UF anticipate the boundary of the legacy network and the proposed design to be?**
Answer: There will be no boundary per-se. We will be swapping out our existing core network for the new one on a box by box basis. If we reduce our box count, we will deploy a new one, and move buildings over one by one. We will be completely removing the existing legacy core network during this migration.

61) **Question: Is the University interested in innovative designs to modernize the deployment, or is the University looking for a hardware update to support the existing feature implementations?**
Answer: We are happy to discuss some design elements; however some elements are committed and being executed on.

62) **Question: How many total interface connections into the core are there currently? Please estimate the number of interfaces and interfaces speeds desired to support connections into the core?**
Answer: The current in-use port count is as follows:

ssrb230a core 6509: 23 1G, 24 10G
csev1 core 6509: 28 1G, 19 10G
mebv1 core 6509: 22 1G, 21 10G
ctx36 core 6509: 17 1G, 23 10G
criv1 core 6509: 20 1G, 24 10G
tig45 core 6509: 17 1G, 21 10G
elmv1 core 6509: 22 1G, 18 10G
aerv1 core 6509: 19 1G, 18 10G
mowv1 core 6509: 15 1G, 10 10G
fifv1 core 6509: 25 1G, 9 10G
smpb151 core 6807: 22 10G
mbig161 core 6807: 22 10G
ecdctr01a core 6506: 7 1G, 12 10G
escv1 core 6506: 7 1G, 12 10G

ssrb230a nexus 6509: 36 10G
ctx36 nexus 6509: 1 1G, 36 10G
escv1 nexus 6506: 12 10G
ecdctr01a nexus 6506: 12 10G