APPENDIX 3

DETAILED VENDOR QUESTIONNAIRE

A All proposed solutions must offer components, which Support the following list of features. Please simply affirm or deny that your proposed solution meets the following requirements by notating each item with "accept", "deny", or "variance". Please explain any variances in your product. This explanation can be inserted into the empty cell under the variance's feature description. The Participant cannot answer "accept," if the feature will not be in production hardware/software by September 1st, 2013. If a feature is not in production hardware/software at the time of the RFP submission, it must be noted in the response.

A1 General IPv4/IPv6 functions and routing protocols:

	Feature description	Response
A1.1	Support IPv4 and IPv6 addresses on routed physical interfaces, sub-interfaces, and virtual interfaces (including loopbacks).	
A1.2	Support IPv4 and IPv6 secondary addresses on routed physical interfaces, sub-interfaces and virtual interfaces.	
A1.3	Support multiple loopback addresses	
A1.4	Be able to route IPv4 packets between physical ports, sub-interface, and virtual interfaces (including loopbacks)	
A1.5	Be able to route IPv6 packets between physical ports, sub-interfaces, and virtual interfaces (including loopbacks)	
A1.6	Support the assignment of IEEE 802.1Q tags to routed sub-interfaces	
A1.7	Support static IPv4 and static IPv6 routes	
A1.8	Support a standards compliant implementation of OSPFv2 for IPv4 described by RFC 2328 that functions on both physical interfaces, sub-interfaces and virtual interfaces	
A1.9	Support a standards compliant implementation of OSPFv3 for IPv6 described by RFC 5340 that functions on both physical interfaces, sub-interfaces and virtual interfaces	
A1.10	Support standards compliant implementation of IS-IS for IPv4 described by RFC 1195, ISO 10589 that functions on both physical interfaces, sub-interfaces and virtual interfaces	
A1.11	Support standards compliant implementation of IS-IS for IPv6 described by RFC 5308 that functions on both physical interfaces, sub-interfaces and virtual interfaces	
A1.12	Support RFC 5303 ISIS – Three-way Handshake for IS-IS Point-to-Point adjacencies	

A1.13	Support RFC 5120 – M-ISIS: Multi Topology (MT) Routing in Intermediate Systems to Intermediate Systems	
A1.14	Support a standards compliant implementation of BPGv4 as described by RFC 4271	
A1.15	Support multi-protocol extensions for BGP (MP-BGP) as described in RFC 4760	
	Have a BGP implementation that Supports 32-bit autonomous system numbers as described in RFC 4893	
A1.16	Have an implementation of PIM-SM that Supports both IPv4 and IPv6 multicast groups and hosts	
A1.17	Support Internet Group Management Protocol version 3 (IGMPv3) as described by RFC 3376 on routed physical interfaces, sub-interfaces and virtual interfaces	
A1.18	Support IGMPv3 for both IPv4 and IPv6	
A1.19	Support re-distribution of routes between routing tables. Directly connected routes, static routes, and routes known by dynamic routing protocols (IS-IS,IS-IS IPv6, OSPFv2, OSPFv3) must be able to be re-distributed into any dynamic routing protocol table, such that a policy can be applied to determine which routes are installed into the target routing table	
A1.20	Support a ping utility to source ICMP echo requests for both IPv4 and IPv6	
A1.21	Support a ping utility that allows specification of the source address as any locally defined IPv4 or IPv6 interface on the device	
A1.22	Support a ping utility that allows specification of the source address as any locally defined IPv4 or IPv6 interface on the device	
A1.23	Support a traceroute utility for both IPv4 and IPv6	
A1.24	Support a traceroute utility that allows specification of the source address as any locally defined IPv4 or IPv6 interface on the device	
A1.25	Support a traceroute utility that allows specification of the source address as any locally defined IPv4 or IPv6 interface on the device	
A1.26	Support IPv4 static ARP entries	
A1.27	Support static IPv6 neighbor table entries	
A1.28	Support of DiffServ code point setting/manipulation QoS and CoS for IPv4 and IPv6	

A1.29	Support a minimum of 100 virtual routing and forwarding (VRF) instances	
A1.30	Support a ping utility that allows specification of the source address as any locally defined IPv4 or IPv6 interface on the device, including addresses assigned in a "VRF-Lite" type context	
A1.31	Support a traceroute utility that allows specification of the source address as any locally defined IPv4 or IPv6 interface on the device, including addresses assigned in a "VRF-Lite" type context	

A2 General Layer 2 protocols and functions:

	Feature description	Response
A2.1	Support bridging functions such that the assignment of a VLAN ID on multiple non-routed (L2) physical ports and a single routed virtual interface comprise a single broadcast domain for the collective set. For example, VLAN 100 could be assigned to four non-routed (L2) physical ports and one routed virtual interface, such that a broadcast packet sourced on the virtual interface would be forwarded to all four physical ports	
A2.2	Fully Support IEEE 802.1Q on non-routed ports	
A2.3	Support broadcast storm control and multicast storm control for all ports operating as L2-only bridging ports	
A2.4	Support jumbo frame sizes of 9,000 bytes on all routed physical interfaces, routed sub-interfaces, routed virtual interfaces, and ports operating as L2-only bridging ports	
A2.5	Support assignment of static L2 (MAC) forwarding table entries for L2-only interfaces	
A2.6	Support of IEEE 802.1p type QoS	
A2.7	Support Q-in-Q VLAN tag stacking	
A2.8	Support VLAN tag rewriting	

A3 Security related:

	Feature description	Response
A3.1	Support the ability to mirror traffic from a set of source ports to single destination port in hardware	
A3.2	Support a minimum of two traffic mirroring sessions	
A3.3	Be able to apply an IPv4 access control list (ACL) in hardware to both inbound and outbound traffic for all routed physical interfaces, sub-interfaces, and virtual interfaces (including loopbacks)	

A3.4	Be able to apply an IPv6 ACL in hardware to both inbound and outbound traffic for all routed physical interfaces, sub-interfaces, and virtual interfaces (including loopbacks)	
A3.5	Support generating flow records with the capability of generating flow records for both IPv4 and IPv6 traffic	
A3.6	Support generating flow records for routed physical interfaces, sub-interfaces, and virtual interfaces	

A4 Management and control plane:

	Feature description	**Response**
A4.1	Support SNMPv2c	
A4.2	Support restricting SNMPv2c access by an access control list or equivalent protection	
A4.3	represent the following information as objects that can be queried via SNMPv2:	
	- Status of physical devices (chassis, fan trays, power supplies, line cards, etc.)	
	- Status of resources (CPU, memory, temperature, etc.)	
	- Status and statistics of internal processes	
	- Status of physical interfaces, sub-interfaces, and virtual interfaces	
	- IPv4 ARP table entries/IPv6 neighbor table entries	
A4.4	Traffic counters for physical interfaces, sub-interfaces, and virtual interfaces	
A4.5	Support remote access to the command line interface (CLI) via SSH	
A4.6	Support restricting router access via SSH by an ACL or equivalent protection	
A4.7	Support access to the CLI via a serial console port	
A4.8	Support encryption of local user passwords stored in the configuration file	
A4.9	Be able to represent its configuration in a text file that can be uploaded to restore a configuration	
A4.10	Support logging to an external syslog server	
A4.11	Log usernames of authorized users to the syslog server at the time of login	

A4.12	Support NTPv3 or NTPv4 (RFCs 1305,5905 respectively)	
A4.13	Support control plane protection such that traffic forwarding, router adjacencies, and CLI management of the device over the serial port continue to function when high volumes of traffic that require handling by the supervisor's CPU are present, including but not limited to ARP storms, broadcast storms, DHCP relay forwarding, and handling traffic with IP options set	
A4.14	OpenFlow support in a hybrid mode	
A4.15	Per VLAN protection in OpenFlow hybrid mode	
A4.16	Support IEEE 802.1AB Link-Layer Discovery Protocol (LLDP)	

A5 Chassis and line card characteristics:

If proposing a non-chassis based solution questions regarding supervisor cards aren't applicable, use NA as a response. You can consider expansion modules and switches in a stack as line cards

	Feature description	**Response**
A5.1	All line cards are hot-swappable	
A5.2	All chassis types Support hot-swapping line cards without degrading the function of other cards in the device	
A5.3	All line cards proposed for 1 Gb operation must Support IEEE 802.3ad link aggregation with other 1 Gb ports on the same and different line cards	
A5.4	All line cards proposed for 10 Gb operation must Support IEEE 802.3ad link aggregation with other 10 Gb ports on the same and different line cards	
A5.5	The chassis proposed must have 40GbE/100GbE capability	

A6 Physical, power and cooling characteristics:

	Feature description	**Response**
A6.1	Be mountable in a standard two-rail 19" communication rack	
A6.2	Support 208V AC	
A6.3	Support N+1 power redundancy power supplies, such that if any one power supply were to fail or lose power, the router would continue to function without any impairment. For example, if a router requires three power supplies, the router must Support four power supplies such that any one power supply could fail and the router would continue to function without impairment	

B1.9	Support IGMPv3 as described by RFC 3376 on routed physical interfaces, sub-interfaces and virtual interfaces in the context of a "VRF-Lite" type instance	
B1.10	Support DHCP relay for IPv4 in a "VRF-Lite" type context	
B1.11	Support assignment of the same VLAN tag ID to one routed sub-interface on every physical port. For example, every physical port on a router could have a routed sub-interface assigned with a VLAN tag of 999 without combining those sub-interfaces into a single VLAN broadcast domain structure	
B1.12	Support Bi-Direction Forwarding Detection (BFD) as described in RFC 5880	

B2 General Layer 2 protocols and functions:

	Feature description	Response
B2.1	Support concurrent use of VLANs 1-4096	
B2.2	Support a minimum of 4,096 hosts in a single layer 2 VLAN construct	

B3 Security related:

	Feature description	Response
B3.1	Support generating flow records for IPv4 and IPv6 traffic on interfaces in a "VRF-Lite" type context	
B3.2	Support GRE tunnels	
B3.3	Support the ability to mirror traffic from a 10 Gb to a 1 Gb port with the understanding that the destination port is 10:1 oversubscribed	

B4 Management and conrol plane:

	Feature description	Response
B4.1	Support SNMPv3	
B4.2	Support restricting SNMPv3 access by an ACL or equivalent protection	
B4.3	Represent the following information as objects that can be queried via SNMP	
	- Transmit and receive light levels on optical transceivers	
	- Serial number and slot position of every field replaceable part	
B4.4	Log the username associated with configuration changes to the syslog server at the time changes are committed	

A6.4	have hot swappable power supplies, such that they are individually removable without impairing or degrading function	

B The following list of features is desirable in the proposed solution. Preference will be given to products that meet the features listed in B. Please simply affirm or deny that your proposed solution meets the following requirements by notating each item with "accept", "deny", or "variance". Please explain any variances in your product. The Participant cannot answer "accept," if the feature will not be in production hardware/software by September 1st 2013. If a feature is not in production hardware/software at the time of the RFP submission, it must be noted.

B1 General IPv4/IPv6 functions and routing protocols:

	Feature description	Response
B1.1	Support assigning individual loopback addresses to each "VRF-Lite" type instance	
B1.2	Support a standards compliant implementation of OSPFv2 for IPv4 as described by RFC 2328 that runs in the context of a "VRF-Lite" type instance	
B1.3	Support a standards compliant implementation of OSPFv3 for IPv6 as described by RFC 5340 that runs in the context of a "VRF-Lite" type instance	
B1.4	Support standards compliant implementation of IS-IS for IPv4 described by RFC 1195 that runs in the context of a "VRF-Lite" type instance	
B1.5	Support standards compliant implementation of IS-IS for IPv6 described by RFC 5308 that runs in a the context of a "VRF-Lite" type instance	
B1.6	Support VRRP version 3 as described in RFC 5798 for both IPv4 and IPv6 on routed physical interfaces, sub-interfaces and virtual interfaces in the context of a "VRF-Lite" type instance	
B1.7	Support PIM-SM as described by RFC 4601 on routed physical interfaces, sub-interfaces and virtual interfaces in the context of a "VRF-Lite" type instance	
B1.8	Support a unicast reverse path forwarding (uRPF) check for both IPv4 and IPv6 unicast traffic on routed physical interfaces, sub-interfaces and virtual interfaces. The uRPF check must discard ingress traffic with a source address that is not within the range of a defined subnet on the receiving routed interface	

B4.5	All optical transceivers should report transmit and receive light levels	
B4.6	All optical transceivers should report transmit and receive light levels	
B4.7	Support some type of management API	

B5 Chassis and line card characteristics:

	Feature description	Response
B5.1	Support "non-stop" forwarding following the failure of the primary supervisor when a redundant supervisor is installed. Failure of the primary supervisor should not result in interruption of service, which includes traffic forwarding and maintaining routing information such that no re-convergence is required by running routing protocols	
B5.2	Support in-service software upgrades when a redundant supervisor is installed	
B5.3	Support replacement of an installed line card with a line card of a different type without requiring a power cycle of the chassis for the new card to function	
B5.4	Support for 3rd party optics	
B5.5	Support for colored CWDM/DWDM optics	
B5.6	SFP+ equipment also Supports SFP optics	
B5.7	Have removable fan trays, such that they are individually removable without impairing or degrading function	
B5.8	Have N+1 redundant fan trays, such that any one fan tray can fail without impairing router function or requiring the router to shut down	

C Metrics/Scaling

Each Participant must answer the following set of questions designed to help identify basic metrics and scaling issues. If more than one type of device is proposed, the values for each type must be specified.

C1	Specify the maximum number of unicast routes Supported. If there is any difference between IPv4 and IPv6, then please specify.	
C2	Specify the maximum number of multicast routes Supported. If there is any difference between IPv4 and IPv6, then please specify.	
C3	Specify the maximum number of unicast routes Supported in a "VRF-Lite type instance". If there is any difference between IPv4 and IPv6, then please specify.	

D3	Specify the maximum power consumption of each type of device.	
D4	Specify the BTU output of each proposed device configuration.	
D5	Specify the maximum BTU output for each type of device.	

E General SDN Strategy

E1	Describe SDN elements in your current product line(s).	
E2	Describe the SND roadmap for your product line(s).	

Signature Date

Printed Name Title

C4	Specify the maximum number of multicast routes Supported in a "VRF-Lite" type instance. If there is any difference between IPv4 and IPv6, then please specify.	
C5	Specify the maximum number of "VRF-Lite" type instances supported.	
C6	Specify all known conditions that result in non-hardware forwarding of IPv4 and IPv6 routed packets.	
C7	Specify the maximum Ethernet frame size Supported. If it is not consistent across all interface types, specify the maximum frame size per interface type.	
C8	Specify the size of the output buffer queue for the interfaces/line cards proposed.	
C9	Describe in detail the control plane protection mechanisms available in each proposed device.	
C10	Specify the per-slot available bandwidth for each proposed chassis.	
C11	Describe the switch fabric bandwidth capacity for each proposed chassis.	
C12	Specify the oversubscription rate for all line cards included in the response.	
C13	Describe any oversubscription that takes place either within a line card or between line cards in your response.	
C14	Describe the line card bandwidth capacity and ASIC arrangement for each proposed line card type.	
C15	List the flavor and versions of flow records you Support.	
C16	Describe your management API if you have one.	
C17	Indicate any line cards in your proposed design that is at full capacity for either 1GbE or 10GbE interfaces given the port counts in figure 2 and figure 3.	
C18	Please describe your 100GbE offering and what we would have to do to add a 100GbE interface to the proposed product.	

D Physical Power and parameters

D1	Specify the number of rack units for each type of chassis in the proposal.	
D2	Specify the power requirements (wattage) of each proposed device configuration when using 208V AC power.	