

---

## Logical overview of NOMX Premium Co-location network for **INET** - setup for the Nordic Markets.

**N.B:** a separate document contains the overview for **Genium INET**

Version 2.1  
February 19<sup>th</sup>, 2011

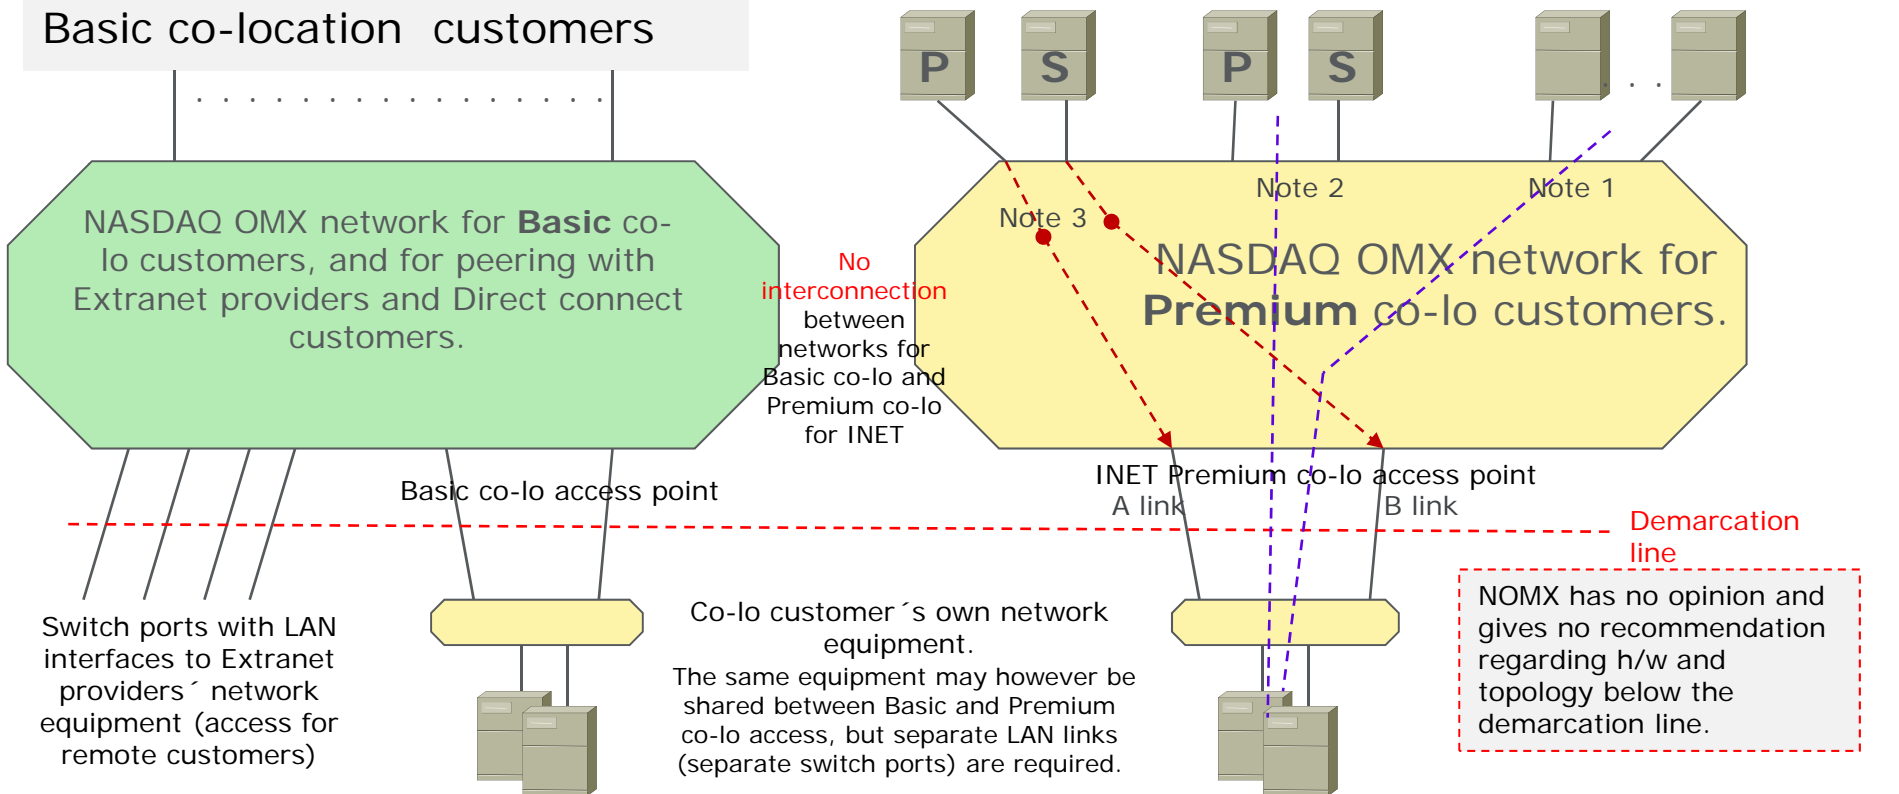
**INET Nordics Production** services in hosts with LAN interface to the Premium Co-lo Service switches.

**All NOMX** services available to both to remote customers and Basic co-location customers

**ITCH IP Multicast** hosts  
Primary & Secondary flows

**ITCH Re-request** hosts;  
Primary & Secondary

**OUCH** hosts; each OUCH account will be set up in two hosts (Primary and Secondary connections).



Notes explained on next page

## Premium Co-lo INET - IP addresses

### IP Unicast traffic

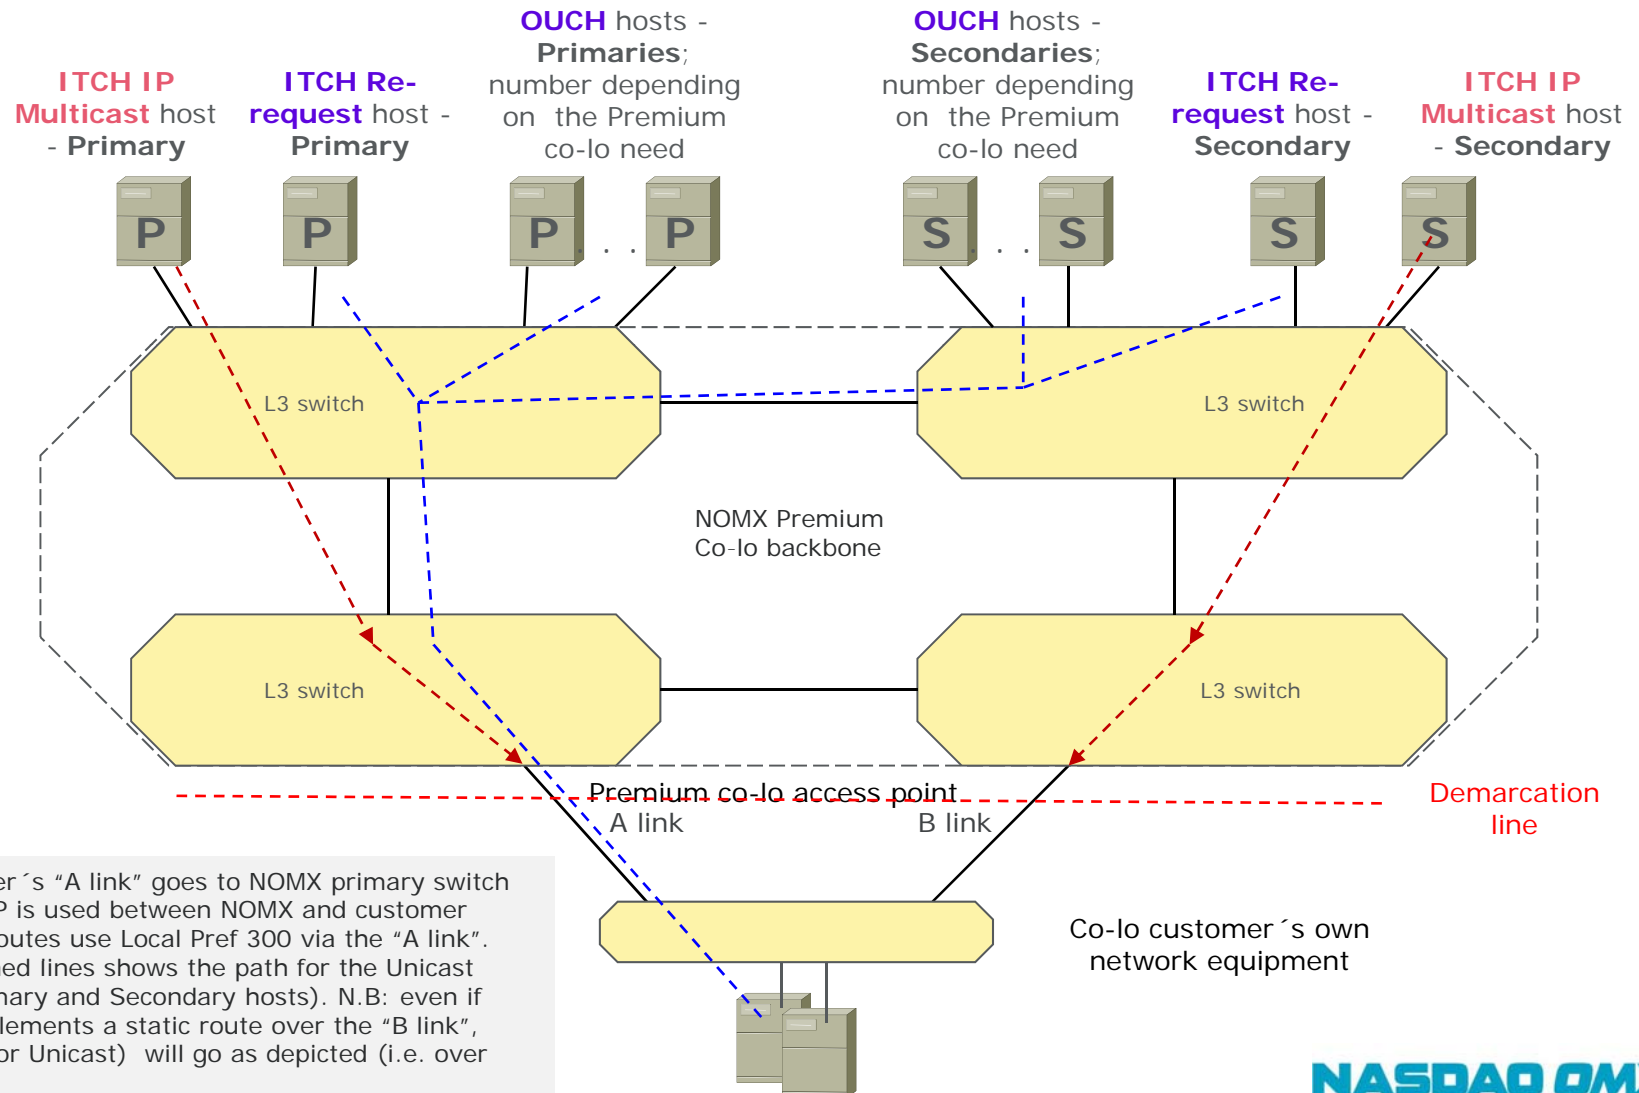
- Note 1: a Premium co-lo customer connects to the OUCH hosts where the OUCH Production accounts are set up. IP net: **193.142.162.0/23**. The destination IP addresses to connect to will be handed out together with each account (i.e. Username & Password for the *SoupTCP* Login Request, sent over the TCP connection established with the TCP port for OUCH). TCP port is 16020.
- Note 2: a Premium co-lo customer using ITCH IP Multicast sends re-request to any of the Re-request hosts. IP net: **193.142.162.0/23** (i.e. same as above). The specific IP addresses and UDP ports as in the layout on Page 11.

### IP Multicast traffic

- Note 3: for a Premium co-lo customer having ordered ITCH Multicast, NOMX will enable the Primary ITCH Multicast Production flows to one of the switch ports (LAN link A) of customer's access point, and the Secondary Multicast flow to the other switch port (LAN link B); and push out the flows. IP Multicast group addresses: **233.74.125.35** (Primary flow) and **233.74.125.34** (Secondary flow). Source IP nets: **192.176.4.160/29** (Primary) and **192.176.4.168/29** (Secondary). The specific source IP addresses and the UDP ports as in the layout on Page 10.

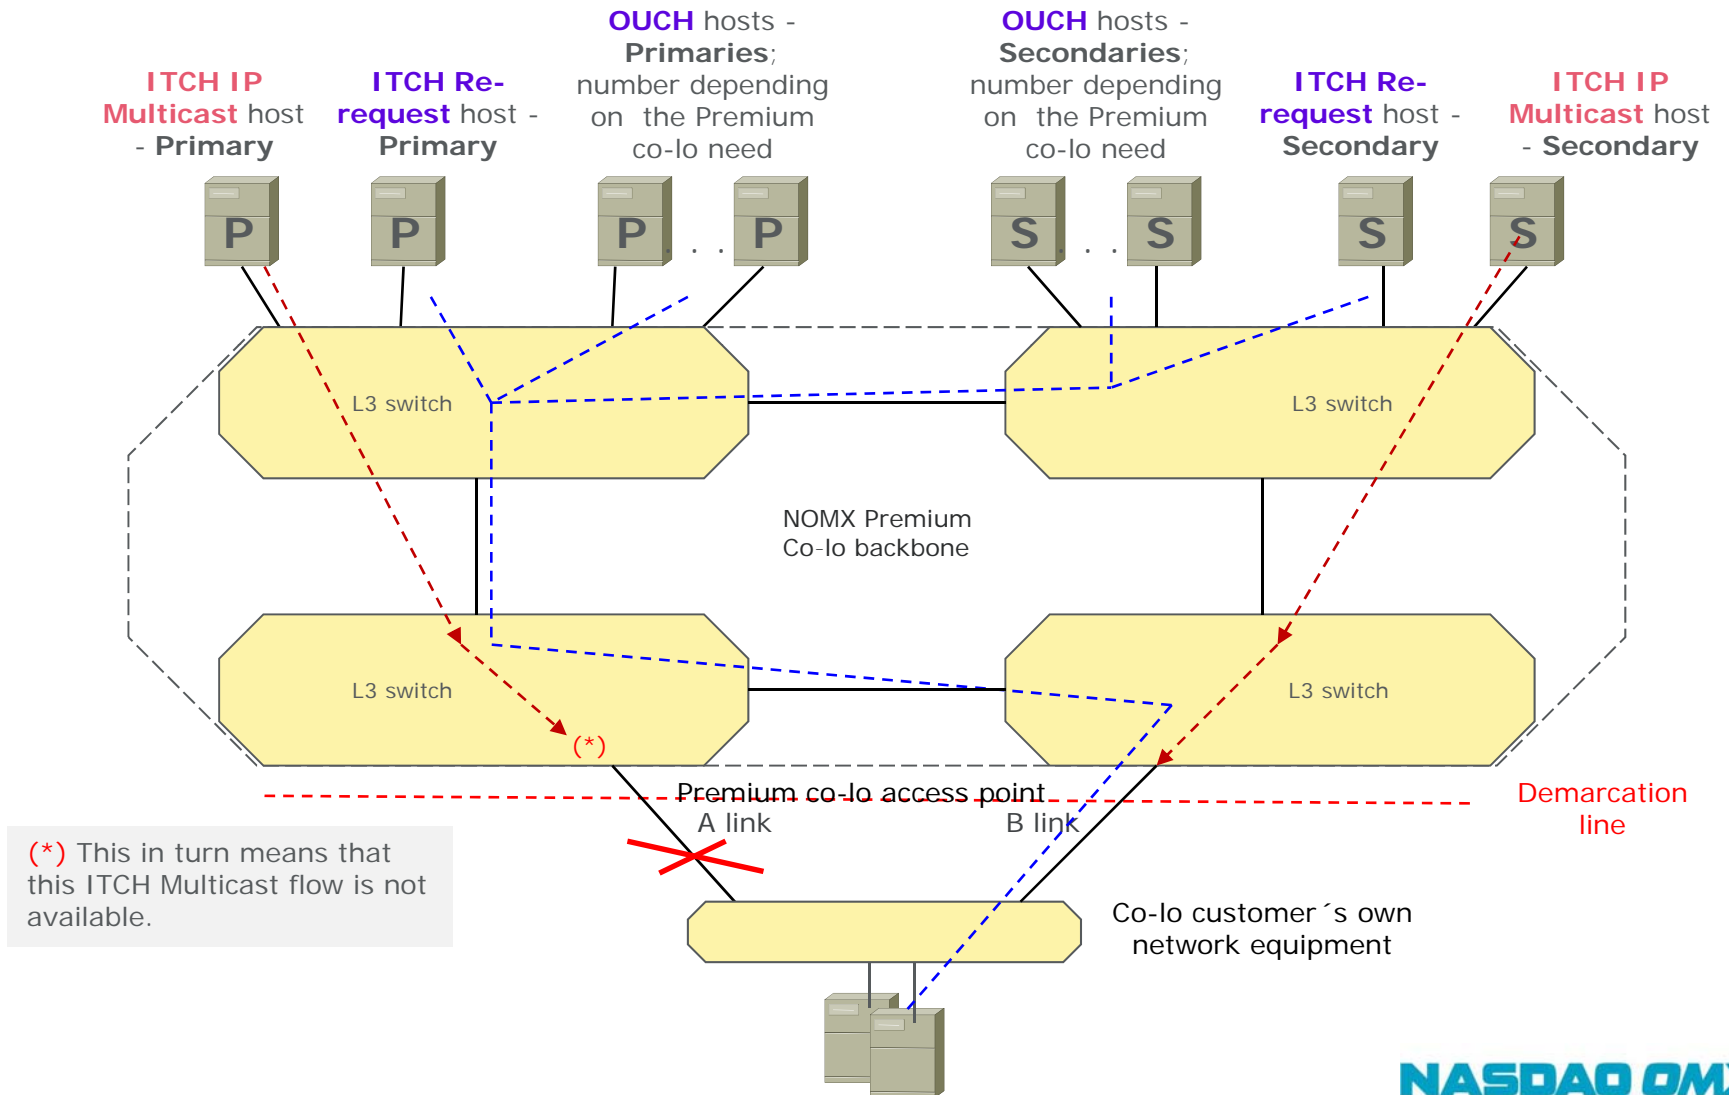
## This is a close up of the Premium Co-lo part from page 2.

All host services are available at any time, but the Co-lo customer needs to be aware that the shortest possible path applies to the Primary hosts, and the services in the Secondary hosts gives additional switch hop.



Co-lo customer's "A link" goes to NOMX primary switch chassi. E-BGP is used between NOMX and customer equipment. Routes use Local Pref 300 via the "A link". The blue dashed lines shows the path for the Unicast traffic (to Primary and Secondary hosts). N.B: even if customer implements a static route over the "B link", traffic back (for Unicast) will go as depicted (i.e. over the "A link").

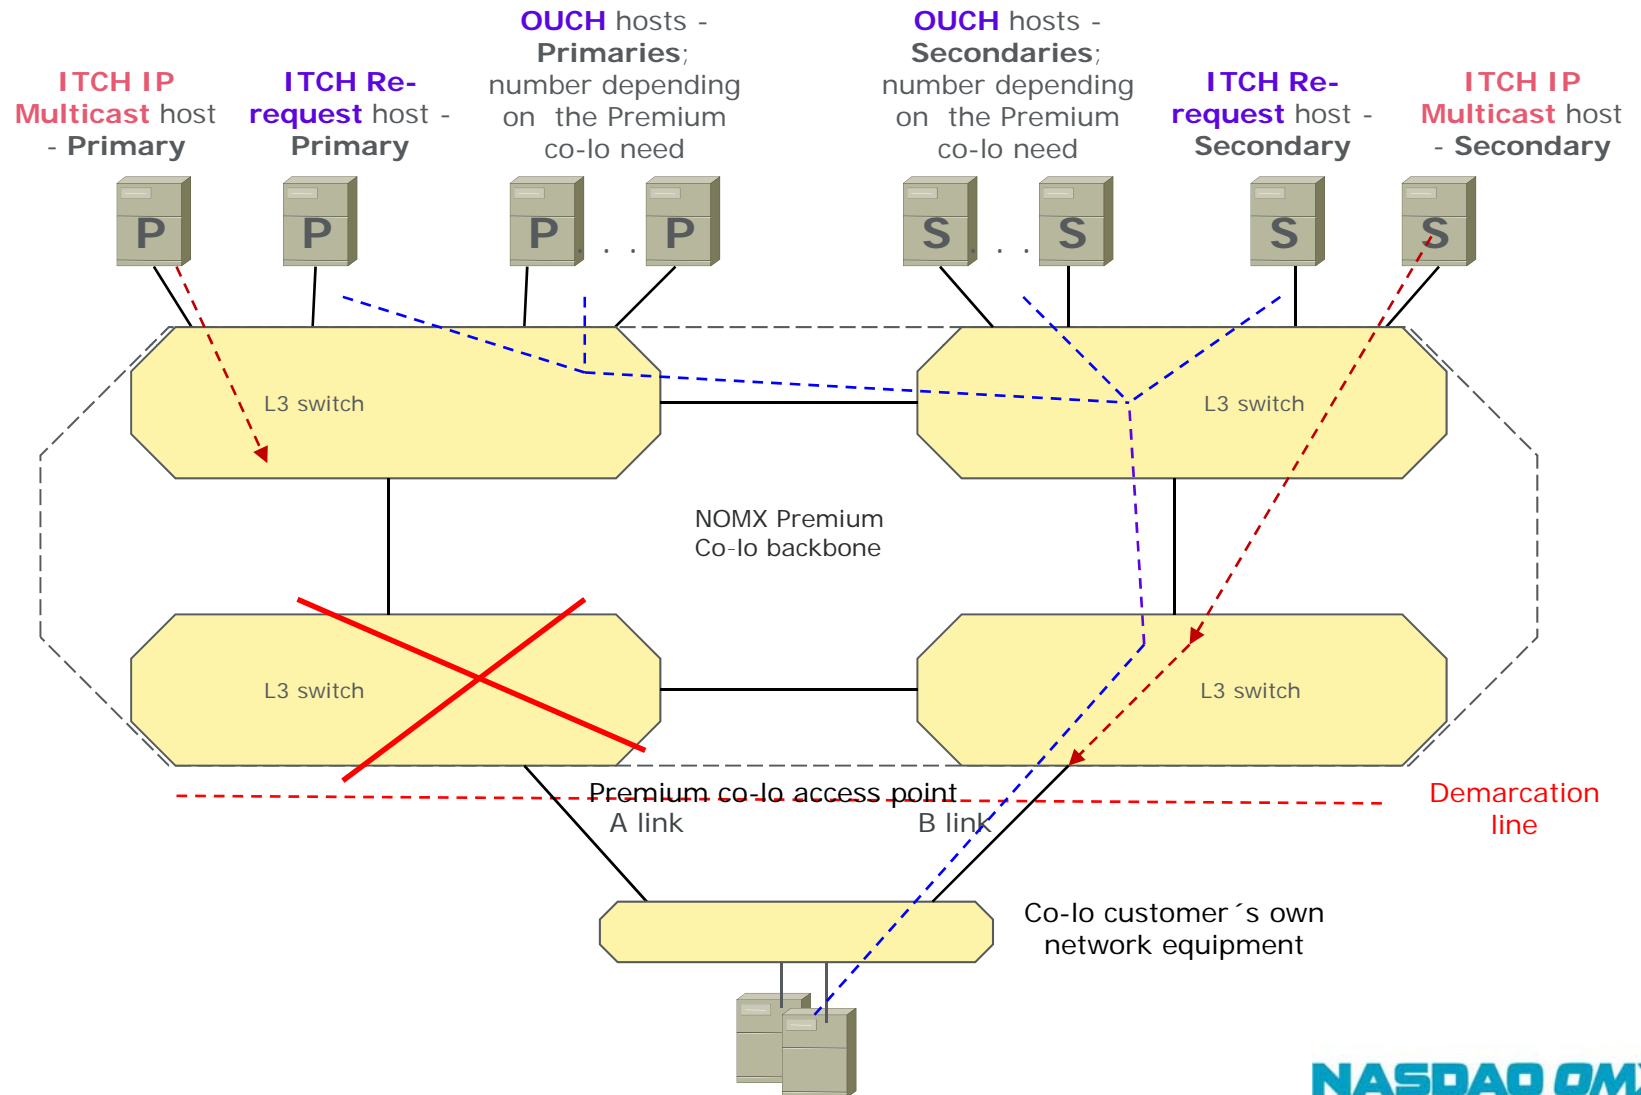
This layout shows the case when customer's "A link" is down.



(\*) This in turn means that this ITCH Multicast flow is not available.

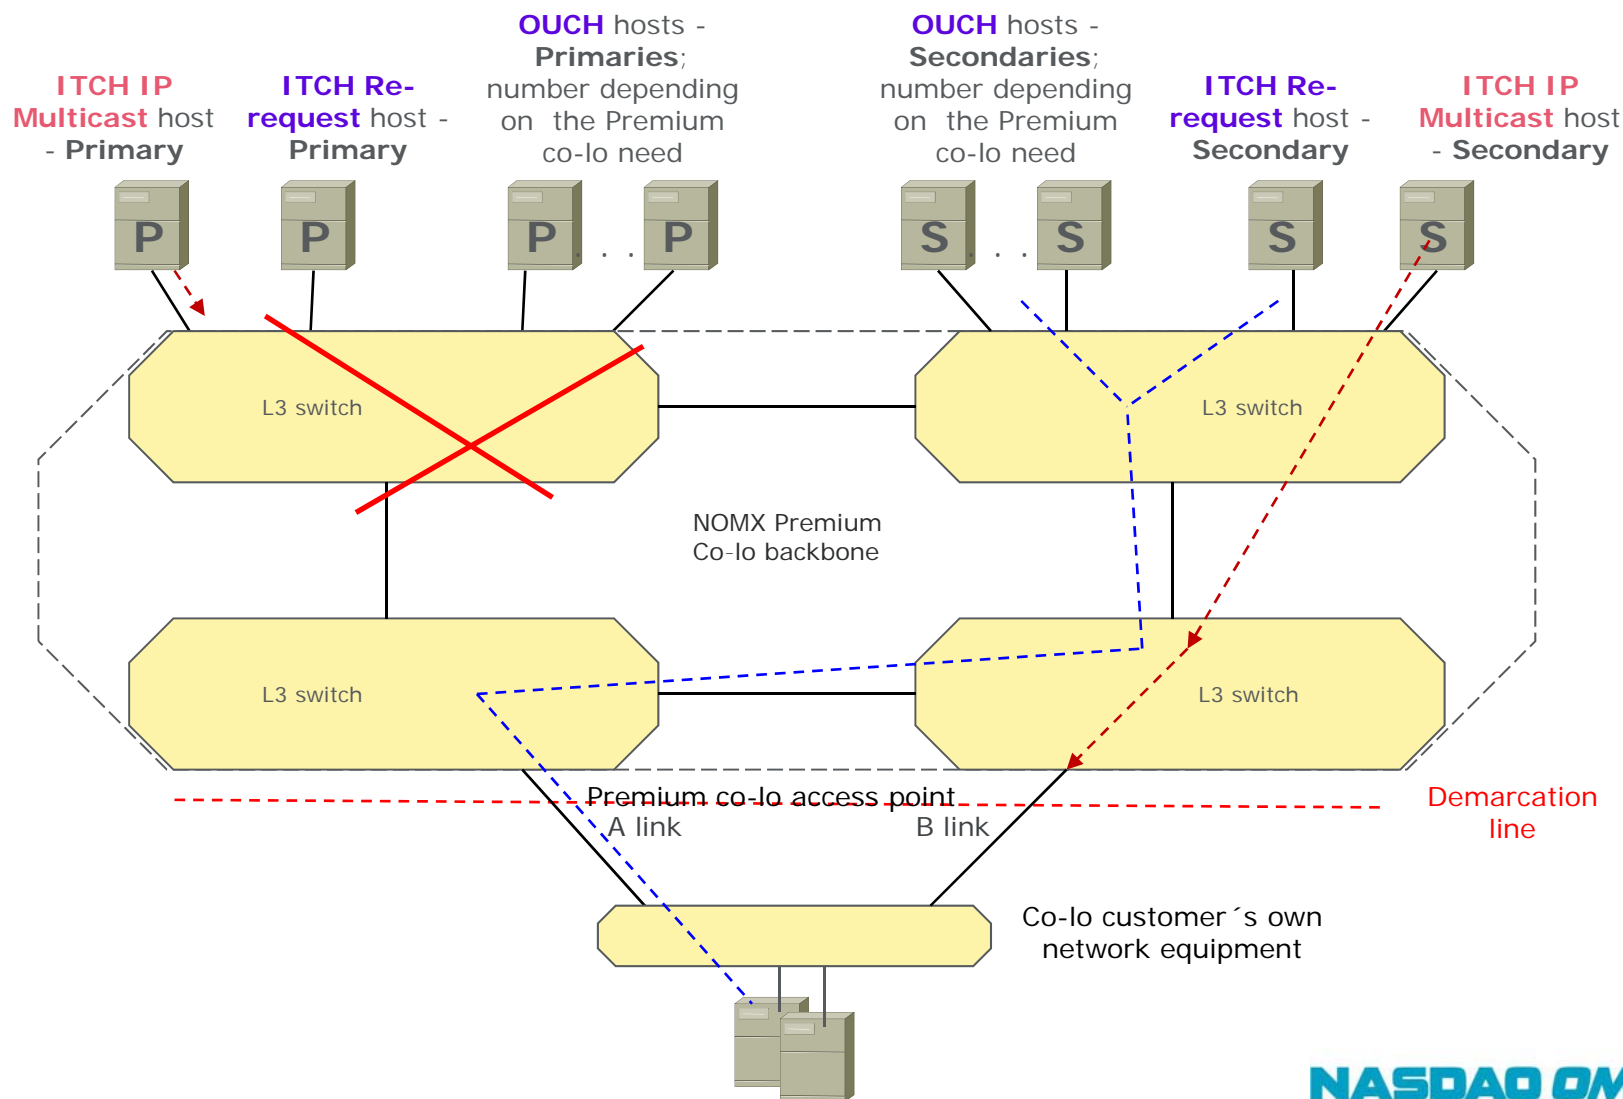
This layout shows the case when the NOMX switch for customer's "A link" is down.

Same availability as on previous page (only ITCH Multicast affected), but different path.

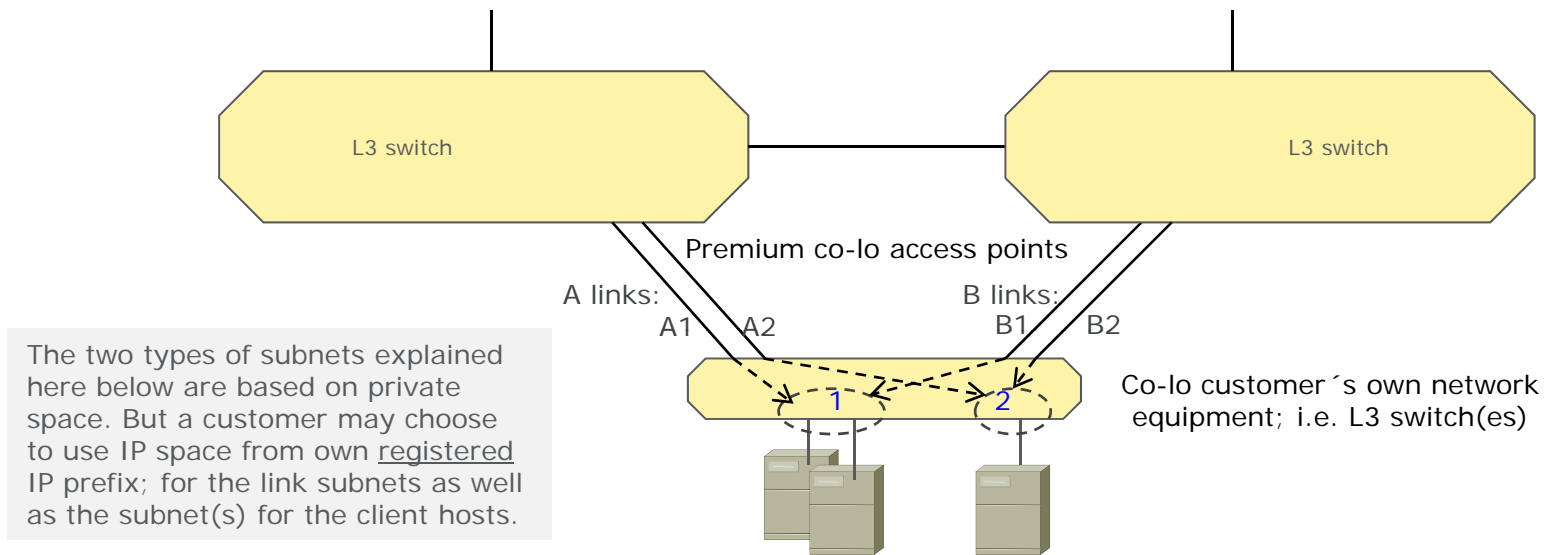


This layout shows the case when the NOMX switch for Primary hosts are down.

Only services in the Secondary hosts are available.



The previous pages have assumed a customer access-point consisting of one pair of access links (i.e. an "A link" and a "B link" consuming one pair of switch ports in the NOMX switches). More access links are possible, but will demand more IP nets to be assigned.



For each individual LAN link a small subnet (30 bits mask) will be assigned from the private space 10.212.18.0/23. In this example with four links (two "A" and two "B"), there will hence be four 10.212.18.x/30 subnets.

Each customer will be assigned an IP subnet (allocated from the private space 10.62.0.0/16) for the hosts acting as INET clients. But if more than one link pair has been ordered, there will also be more than one subnet assigned for the hosts. As depicted for this example, two subnets are assigned. The customer announces "client subnet 1" only over the "A1,B1" pair and "client subnet 2" only over the "A2,B2" pair. And NOMX will route traffic accordingly, where the arrows in the picture shows over which link pair NOMX may route to "client subnet 1" and "client subnet 2" respectively. See however here below as the need of subnets is reduced depending on ITCH Multicast setup.

A customer may request that NOMX pushes out ITCH Multicast on only one link pair. The customer can thus dedicate a link pair for ITCH Multicast only. If say, "A2,B2" is only used for ITCH Multicast, there is no traffic to clients for NOMX to route via "A2,B2". This in turn means that "client subnet 2" is not needed and all customer's hosts can reside on "client subnet 1".

Even though INET is not accessed via the same NOMX switches as Genium INET, the "client subnet(s)" mentioned here can be the same. Hence, the customer can choose whether different "client subnets" shall be assigned or the same subnet used for both INET and Genium INET. It is also possible to use the same subnet for Basic co-lo.

---

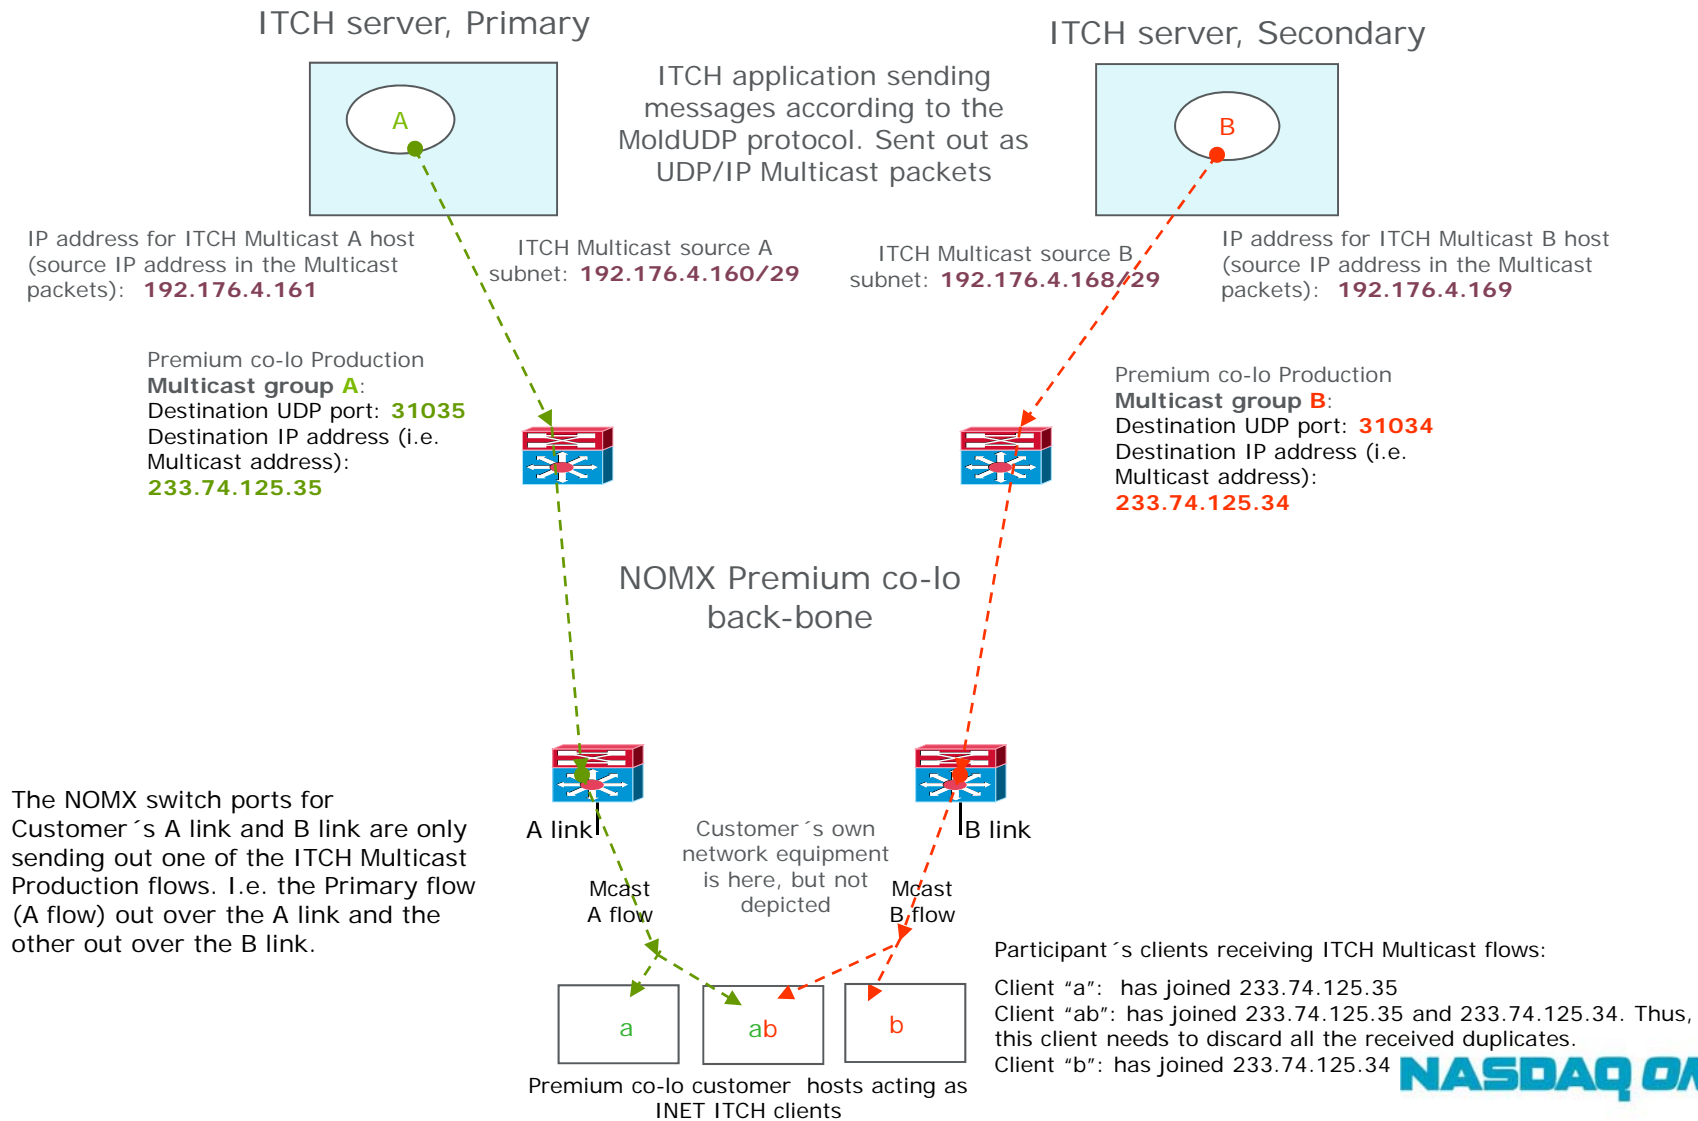
For those not familiar with the INET IP Multicast principles, the following description is found on the member web, path *INET Nordic -> Connectivity and Protocols* (and under headline Connectivity Guidelines): [ITCH UDP Multicast Offering](#).

The layouts in the above document do however not explain the setup from a co-lo perspective, but all the IP addresses and UDP ports apply to those using the Basic co-lo service. For the Premium co-lo setup, the IP addresses (and UDP ports) are different. The following pages contain layouts with details for ITCH Multicast (and Re-request) Production flows applicable to a Premium co-lo customer. The layouts are with same type of ITCH client examples as the above referenced document.

**Premium co-lo ITCH flow based on IP Multicast –  
Production UDP ports and IP addresses**

A flow

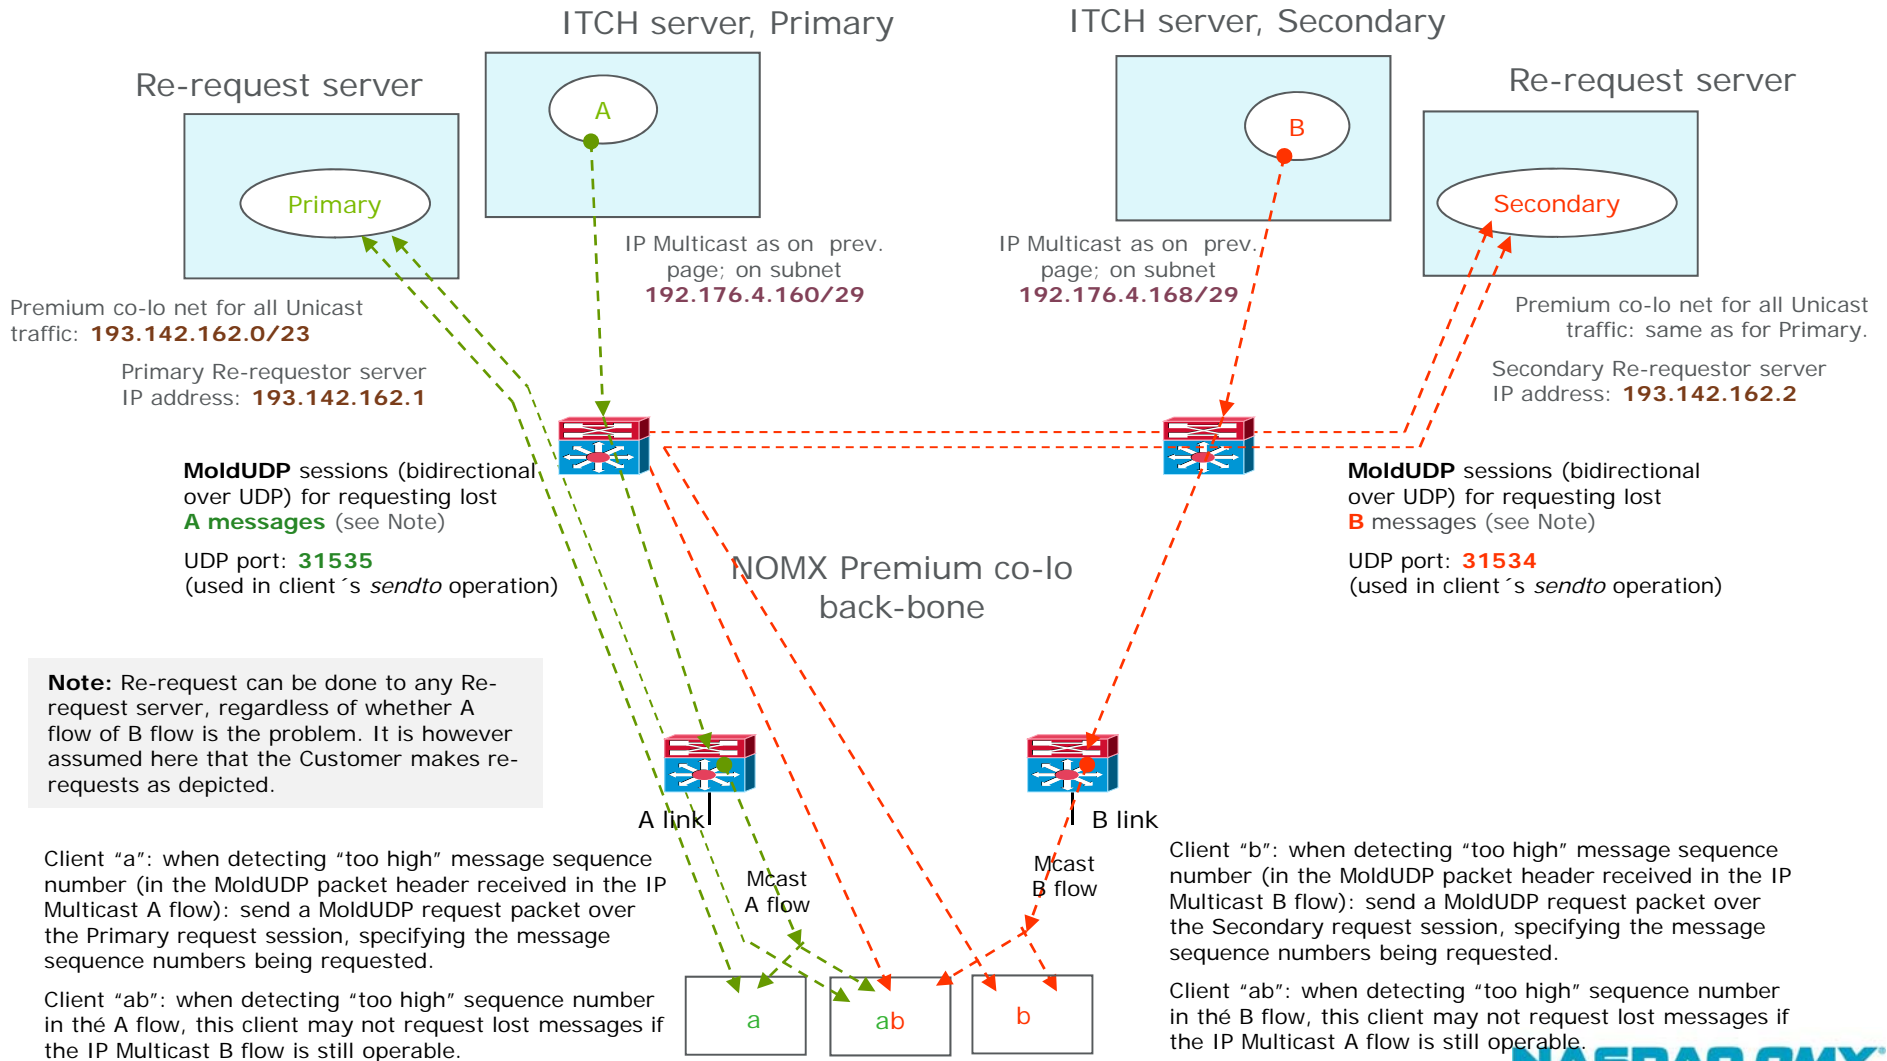
B flow



Premium co-lo ITCH Re-request MoldUDP sessions (for requesting lost messages in the IP Multicast flow) –  
**Production** UDP ports and IP addresses

Primary

Secondary



**Note:** Re-request can be done to any Re-request server, regardless of whether A flow of B flow is the problem. It is however assumed here that the Customer makes re-requests as depicted.

Client "a": when detecting "too high" message sequence number (in the MoldUDP packet header received in the IP Multicast A flow): send a MoldUDP request packet over the Primary request session, specifying the message sequence numbers being requested.

Client "ab": when detecting "too high" sequence number in the A flow, this client may not request lost messages if the IP Multicast B flow is still operable.

Client "b": when detecting "too high" message sequence number (in the MoldUDP packet header received in the IP Multicast B flow): send a MoldUDP request packet over the Secondary request session, specifying the message sequence numbers being requested.

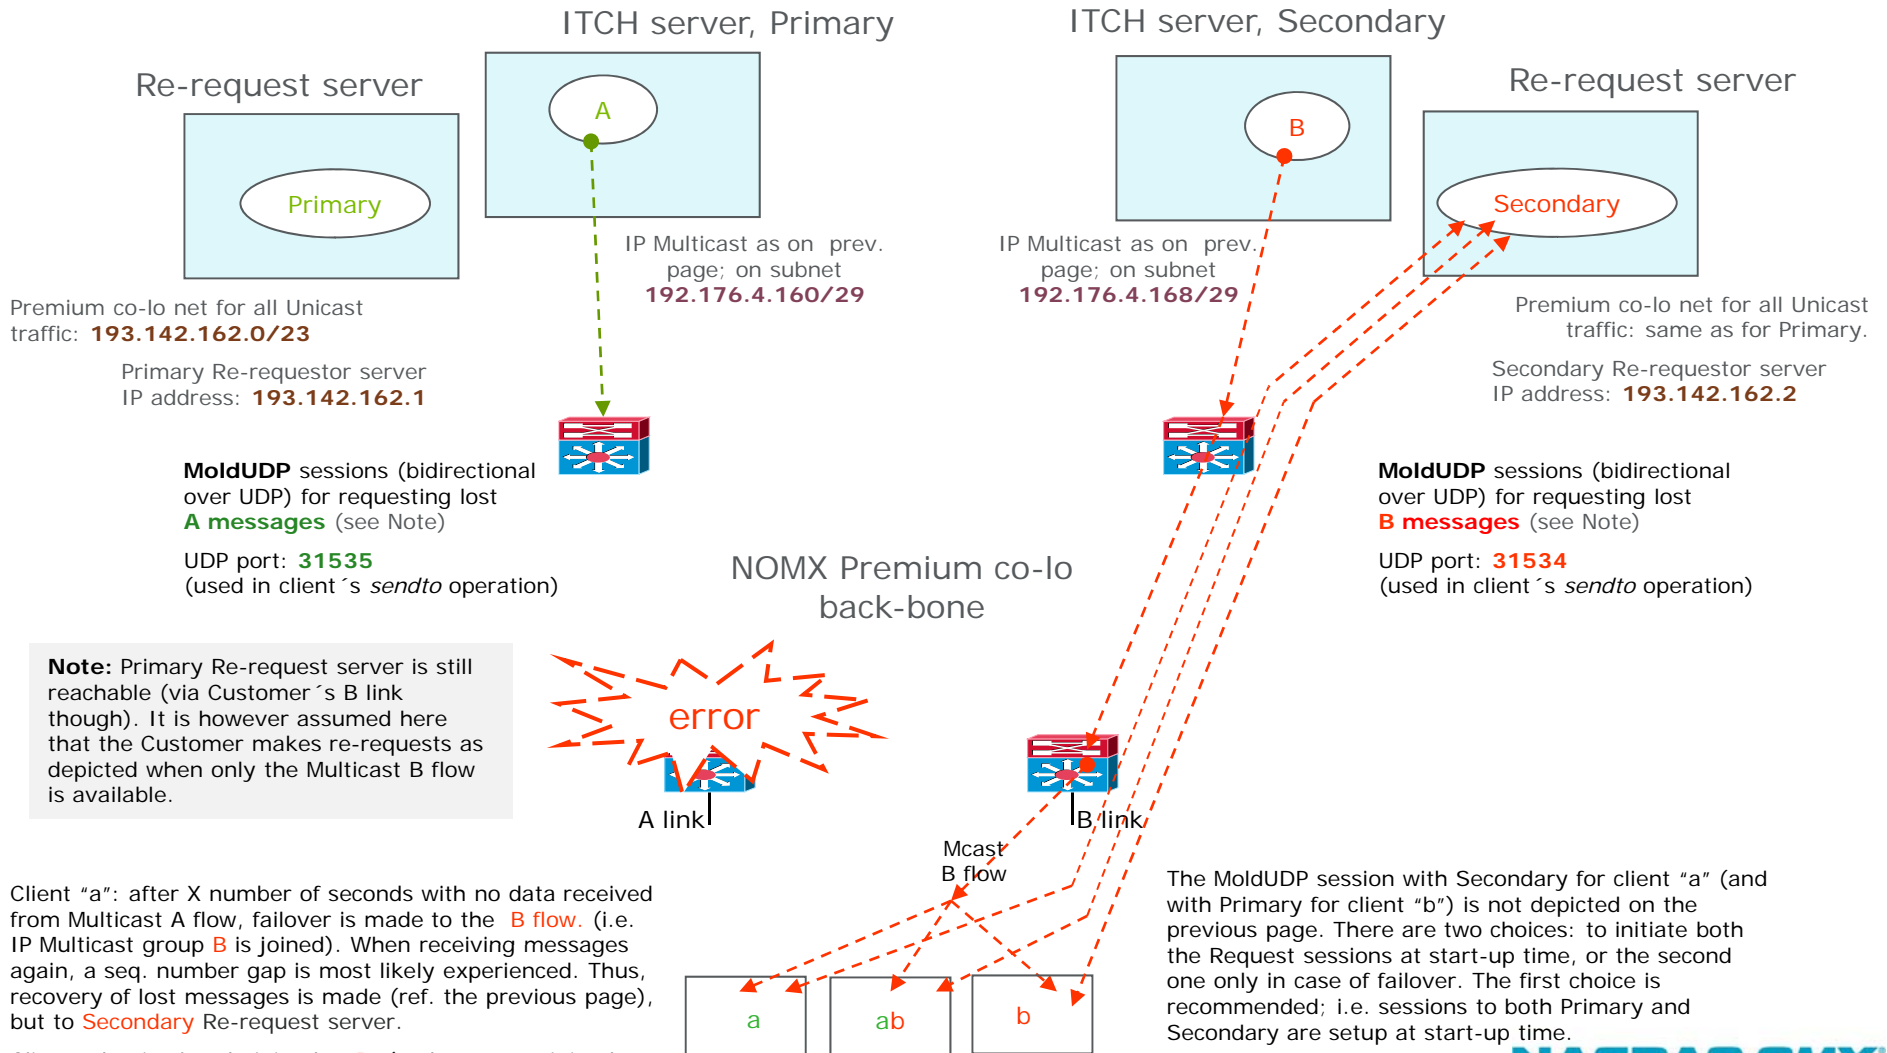
Client "ab": when detecting "too high" sequence number in the B flow, this client may not request lost messages if the IP Multicast A flow is still operable.



Premium co-lo ITCH Multicast for  
Production - Failover case I

Primary

Secondary



**Note:** Primary Re-request server is still reachable (via Customer's B link though). It is however assumed here that the Customer makes re-requests as depicted when only the Multicast B flow is available.

Client "a": after X number of seconds with no data received from Multicast A flow, failover is made to the **B flow**. (i.e. IP Multicast group B is joined). When receiving messages again, a seq. number gap is most likely experienced. Thus, recovery of lost messages is made (ref. the previous page), but to **Secondary** Re-request server.

Client "ab": is already joined to B (and may stay joined to A).

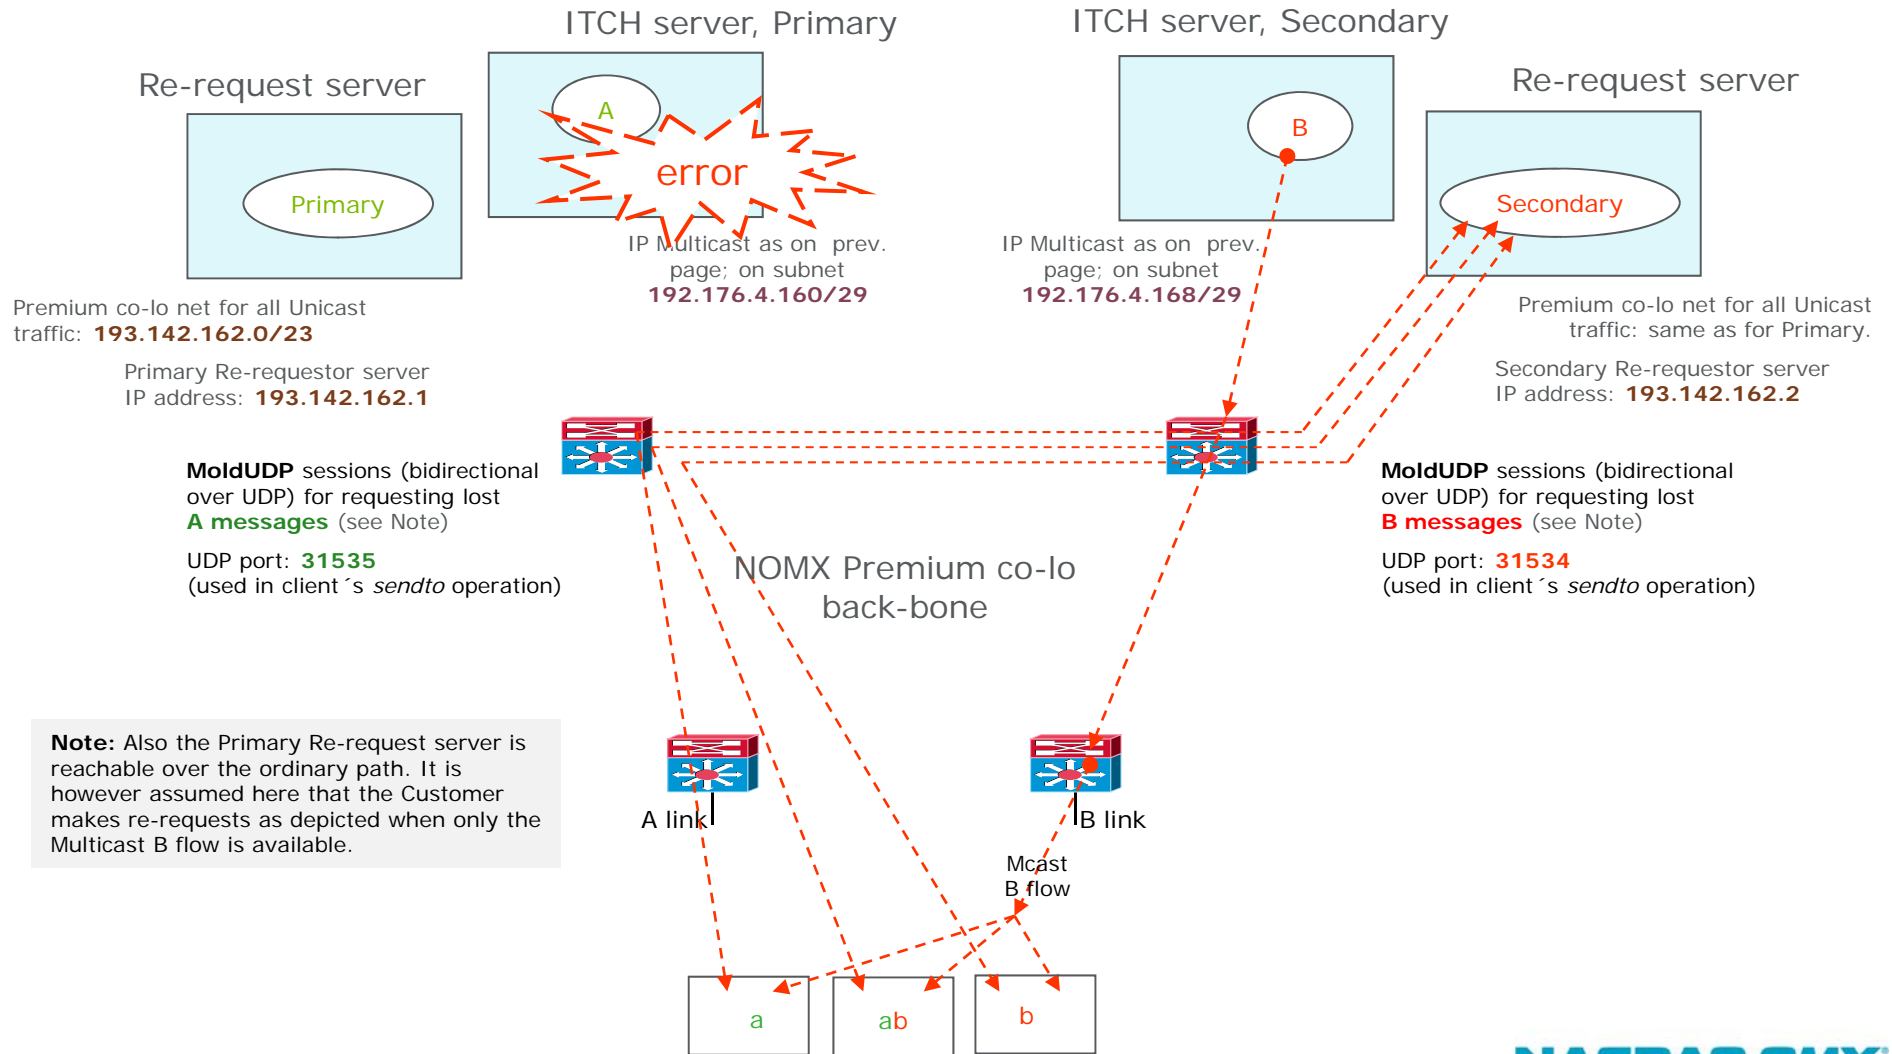
The MoldUDP session with Secondary for client "a" (and with Primary for client "b") is not depicted on the previous page. There are two choices: to initiate both the Request sessions at start-up time, or the second one only in case of failover. The first choice is recommended; i.e. sessions to both Primary and Secondary are setup at start-up time.



Premium co-lo ITCH Multicast for  
Production - Failover case II

Primary

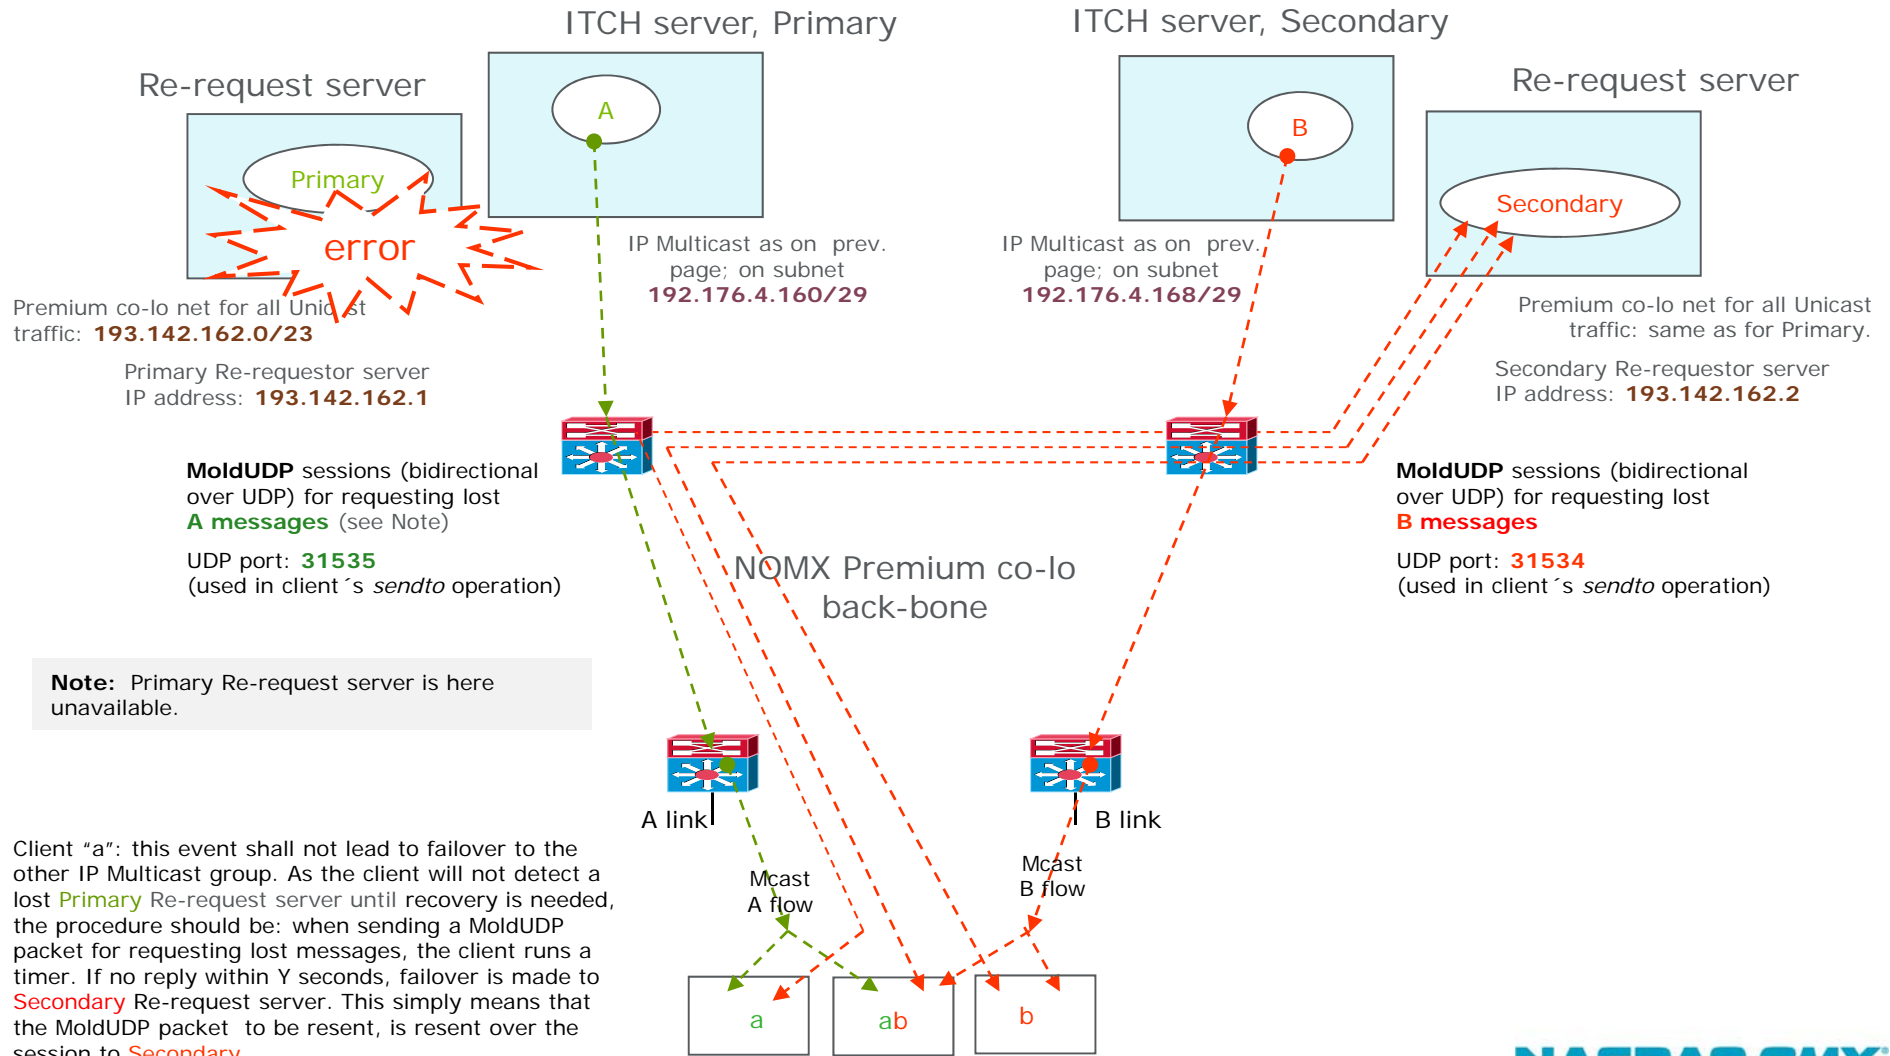
Secondary



## Premium co-lo ITCH Multicast for Production - Failover case III

# Primary

# Secondary



**Note:** Primary Re-request server is here unavailable.

Client "a": this event shall not lead to failover to the other IP Multicast group. As the client will not detect a lost **Primary** Re-request server until recovery is needed, the procedure should be: when sending a MoldUDP packet for requesting lost messages, the client runs a timer. If no reply within Y seconds, failover is made to **Secondary** Re-request server. This simply means that the MoldUDP packet to be resent, is resent over the session to **Secondary**.