Overview of AMS
(CCSDS Asynchronous Message Service)

DARPA DTN
Phase 2 Kickoff

Arlington, VA

9 August 2006

Scott Burleigh
Systems Engineering Section
Jet Propulsion Laboratory, California Institute of Technology
818.393.3353
Scott.Burleigh@jpl.nasa.gov

Objectives

- Provide a general-purpose system for exchange of short messages in deep space mission operations.
  - Telemetry
  - Commands
- Be simple to use.
  - Minimize applications development cost
  - Minimize software configuration and management cost
- Be ubiquitous.
  - Portable to all mission environments
  - Usable over all data transport systems
  - Suitable for all message exchange operations
  - Scalable to large applications
- Be robust.
  - No single point of failure
  - Tolerant of delay
  - Tolerant of data transport disruption

Key Features

- Core *message bus* model
  - Publish/subscribe by message *subject*.
  - Each application software node subscribes to (and consumes) the information it needs, and publishes the information it produces, without knowing which other modules are currently running.
- Other communication models supported as needed:
  - *Explicit awareness of other nodes*
  - *Private message transmission* to specific nodes
    - E.g., *replies* to published messages
  - *Synchronous (client/server)* communication
  - *"Announcement"* of data to multiple anonymous nodes
- Remote AMS (*RAMS*)
  - Aggregates message publication to minimize bandwidth consumption on constrained links
  - Designed to enable dynamic publish/subscribe over interplanetary distances
  - Generalizes to – in effect – scalable reliable multicast

A single AMS continuum

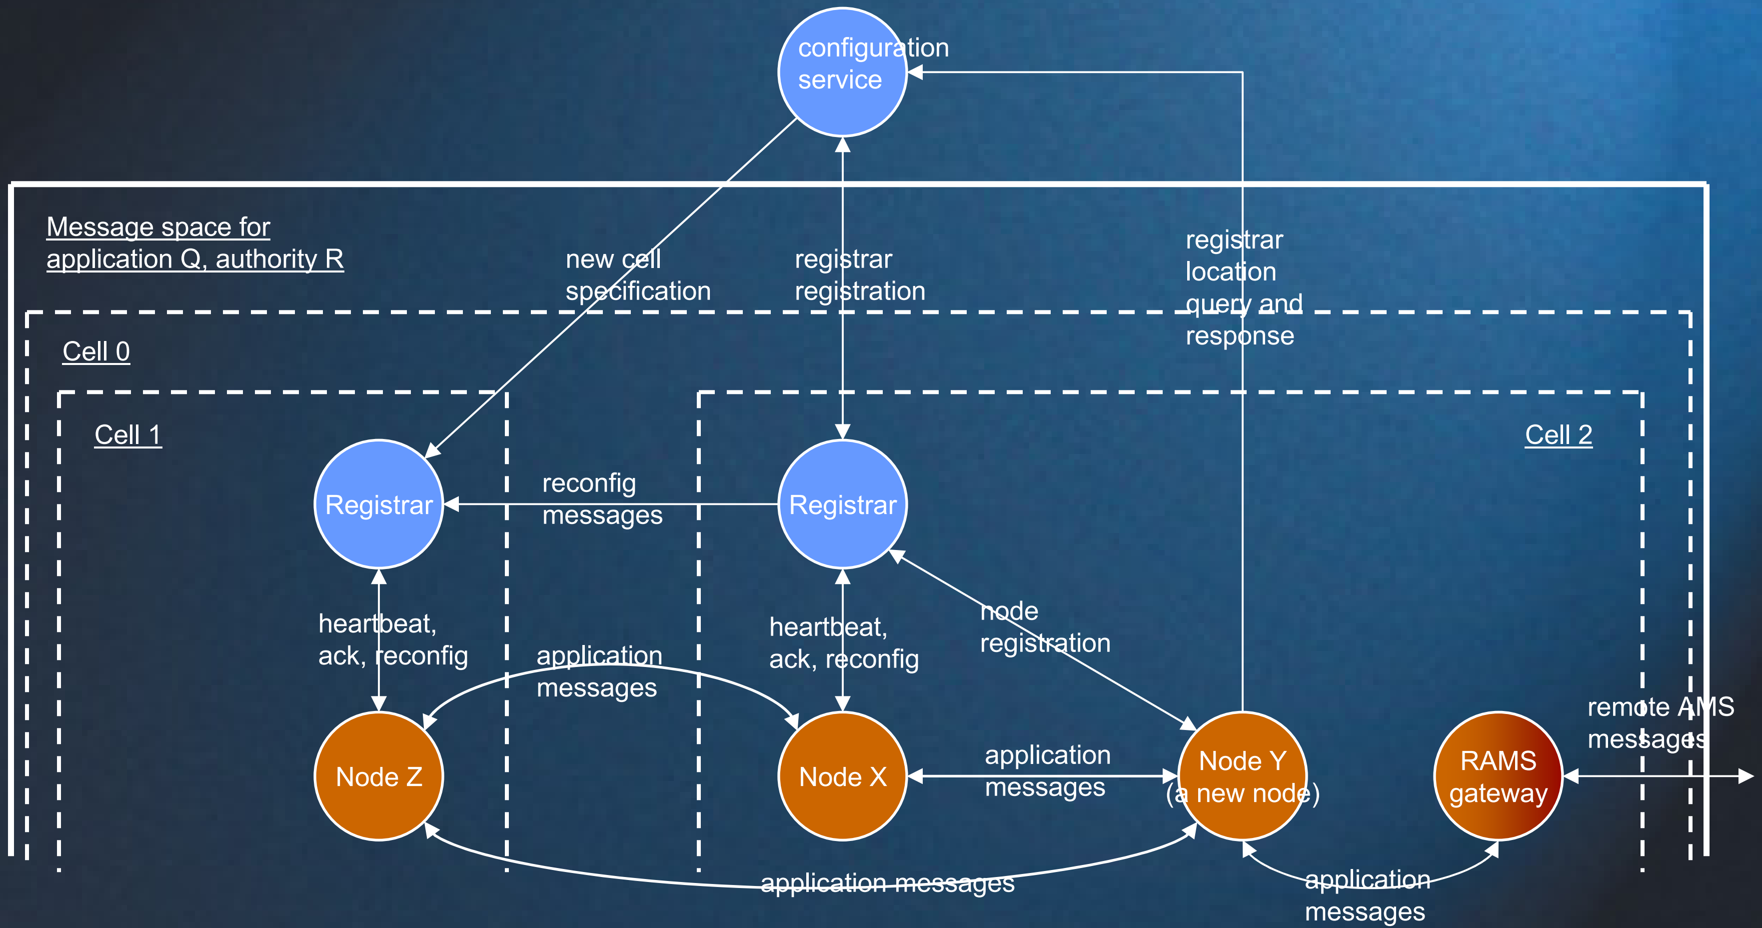

The AMS Protocol Suite

- Meta-AMS
  - Discovery, self-configuration (including subscriptions and subscription cancellations), fault detection, failover, recovery.
  - Messages are exchanged between nodes and configuration servers, between nodes and registrars, between configuration servers, between registrars, and between registrars and configuration servers.
- AMS
  - Application data transmission, incl. queries, replies, announcements.
  - Messages are exchanged directly between nodes (including RAMS gateways, which function as AMS nodes).
- Remote AMS
  - Assertions and cancellations of “petitions”.
    - Aggregated application data transmission.
  - Messages are exchanged between RAMS gateways via DTN.

A multi-continuum "venture"

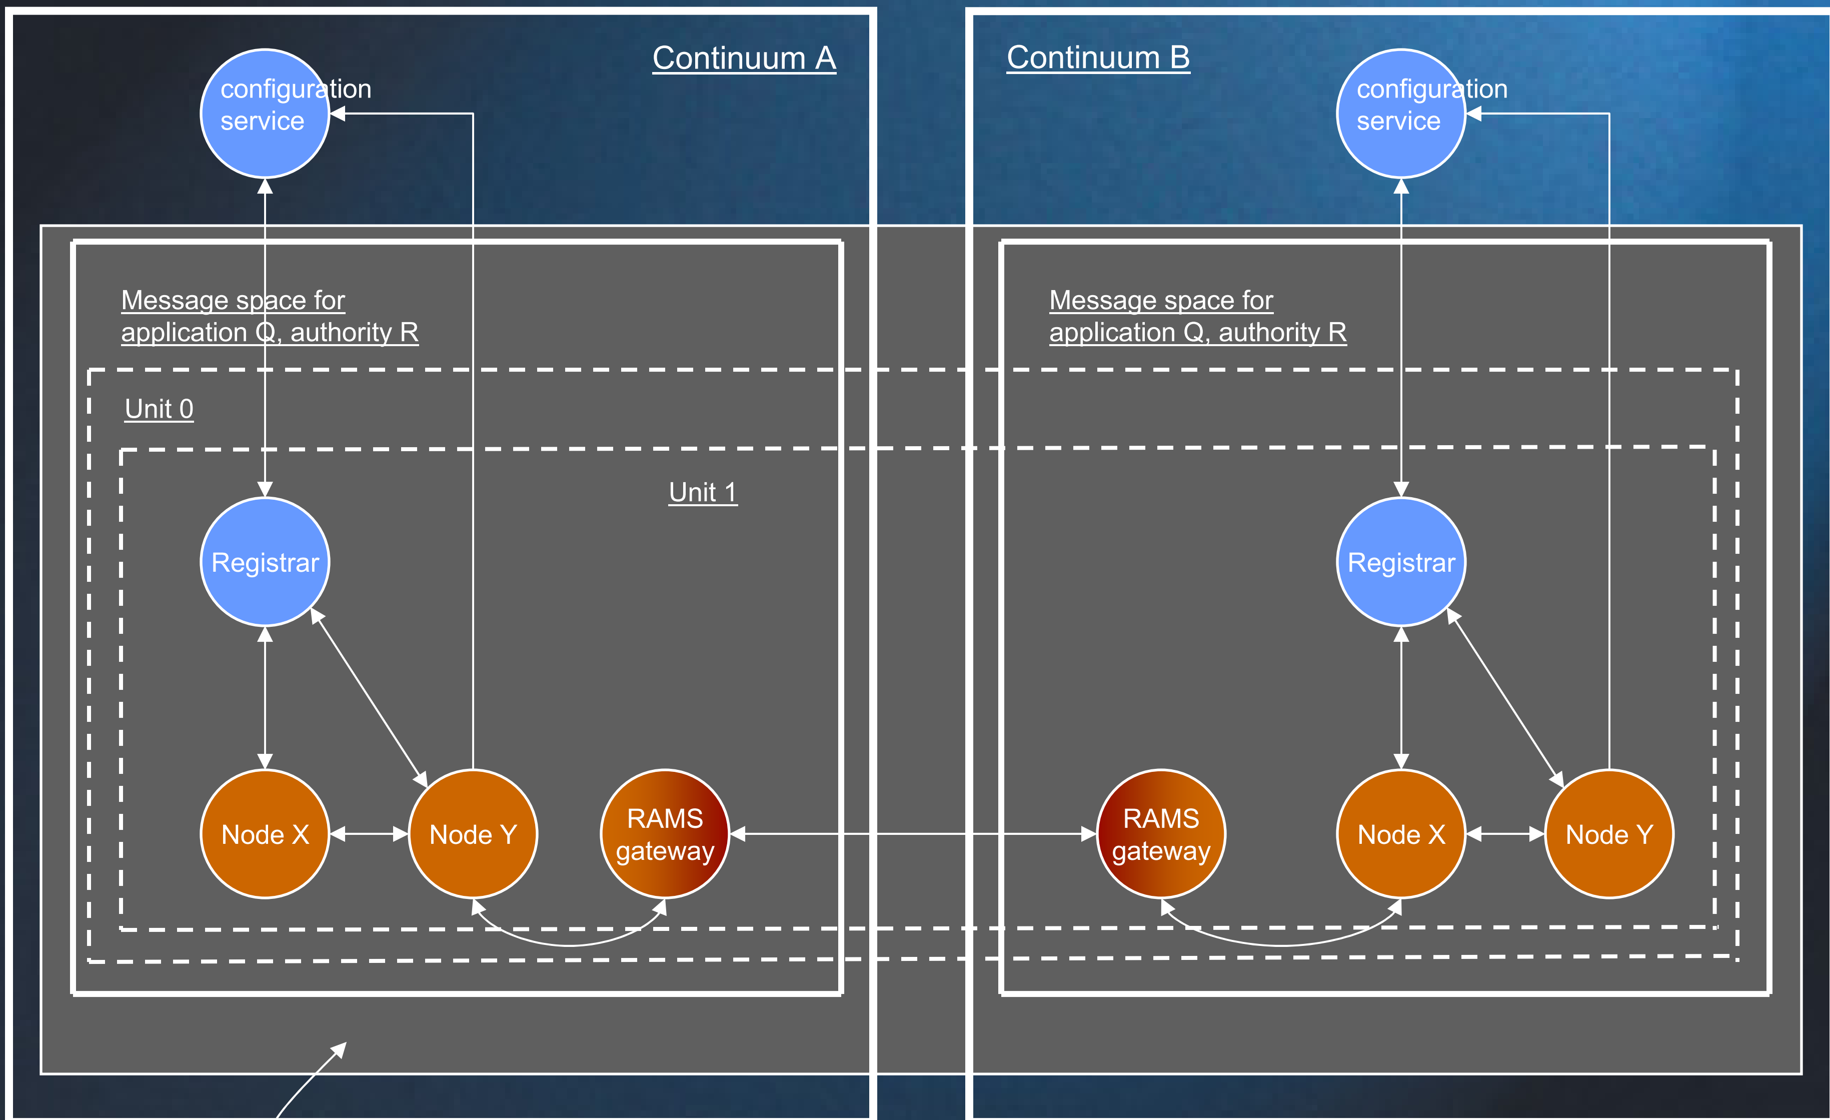

Concatenated application instance ("venture") for application Q, authority R.

Constraining transmissions

- Transmission constraints can be specified in subscriptions (selecting publishers) and in announcements (selecting recipients).
  - Organizational constraint: all – and only – nodes registered in a specified *unit* or in any unit that's wholly contained within the specified unit.
  - Functional constraint: all – and only – nodes declared at registration to be performing a specified *role* in the application.
  - Topological constraint: all – and only – nodes operating within a specified *continuum*.
- Fine-grained control over message publication enables a balance to be struck between latency and bandwidth utilization.
  - Information is pushed rather than pulled, so there is no query/response round trip delay.
  - But information need not ever be pushed to nodes that don't want it.

Security

- Access control
  - List of authorized recipients of messages on a given subject
  - List of authorized producers of messages on a given subject
  - Registration permitted only in authorized application roles
- Authentication
  - Asymmetric encryption
  - Assures authenticity of configuration servers and registrars
  - Basis for access control
- Symmetric encryption of message content

Fault Tolerance

- Preventive maintenance
  - Optional periodic re-issuance of MAMS messages
- Inference of remote node failures
  - Reciprocal heartbeat exchange
- Configuration server failover
- Autonomous recovery

Performance of Reference Implementation

Number of messages sent	*Size of each message (bytes)*	*Messages exchanged per second*	*Data rate (Megabits/sec)*
10,000	**20,000**	**5,337**	**814**
100,000	**2,000**	**25,739**	**393**
1 million	**200**	**107,910**	**165**
10 million	**20**	**154,335**	**23**

Highly preliminary performance measurements, from JPL's Protocol Test Laboratory. Message exchange between a single publisher and a single subscriber on a Gigabit Ethernet. Each node was hosted on a dual-core 3Ghz Pentium-4 running Fedora Core 3. (Don't expect this kind of performance in normal operations!)

AMS vs Multicast (1)

- Both multicast and publish/subscribe result in delivery of a message to each receiver.
- In non-multicast-based publish/subscribe each such message is issued separately by either the publisher or a message broker.
  - So each such message is separately forwarded through the network.
  - Heavy load on network.
- In multicast, each such message is issued by the multicast router that is adjacent to the receiver.
  - The multicast sender sends only to the multicast routers that are adjacent to it.
  - Each such router (until the one at the edge of the transmission) likewise sends only to the adjacent downstream multicast routers.
  - Each multicast router that is adjacent to receivers sends the message just once on its LAN. All receivers on the LAN with sufficient available buffer space acquire the message.
  - So the number of inter-router transmissions can be much smaller.

AMS vs Multicast (2)

- Message forwarding in a RAMS tree operates like multicast.
  - Only the RAMS gateway in each destination continuum sends to the ultimate message receivers, and only to those receivers that are in its own continuum.
  - Each RAMS gateway sends only to its neighboring gateways, acting like a multicast router. Multicast-like load on network.
- Both multicast and RAMS rely on some means of constructing the forwarding tree. For static systems this can be done manually, by static routing. For dynamic systems:
  - For multicast:
    - A different tree must be constructed for each multicast group.
    - Each tree must be dynamically managed as its group's membership fluctuates.
    - So multicast requires a multicast routing protocol.
  - For RAMS:
    - The functional equivalent of the multicast group is the subscription.
    - AMS manages dynamic subscription relationships itself, via MAMS and RAMS.
    - So a single, static forwarding tree can support any number of subscription relationships.
    - No additional routing protocol needed.

RAMS testing exercise

- Objective: demonstrate AMS message exchange over a wide-area network (the Internet), characterize performance.
- Method:
  - Operate AMS continua at JPL, APL, Marshall SFC.
  - Use RAMS to link the three continua.
- This was an artificial use case.
  - Delay over the Internet is low enough to enable all three centers to be in a single continuum. Not what we wanted.
  - To make RAMS necessary, did no firewall modification at any center.
- All traffic had to travel through a third-party routing point at Cal Poly San Luis Obispo – a star-shaped overlay network.
- For routing through this overlay network, used JPL's DTN Bundle Protocol implementation ("ION").

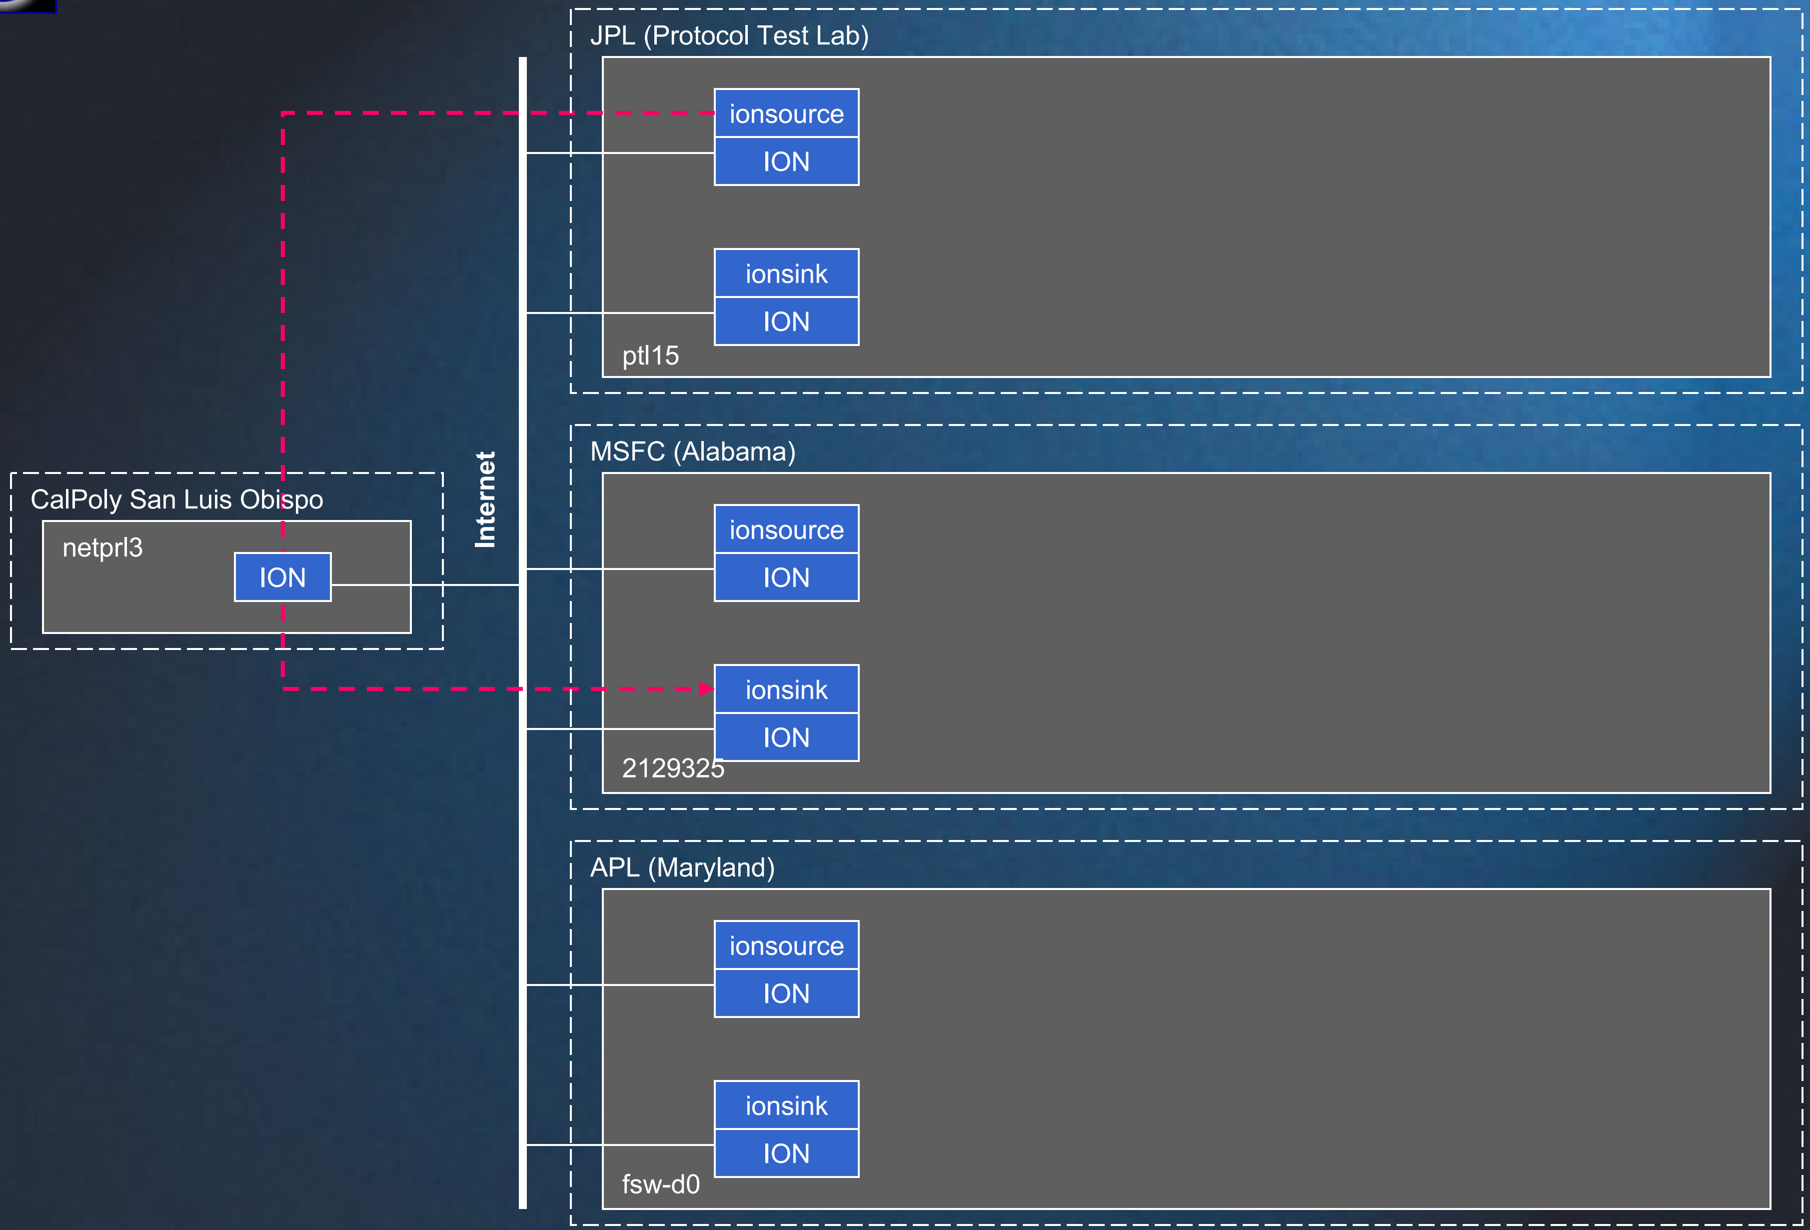
JPL (Protocol Test Lab)
ionsource
ION
ionsink
ION
ptl15
MSFC (Alabama)
ionsource
ION
ionsink
ION
2129325
APL (Maryland)
ionsource
ION
ionsink
ION
fsw-d0
CalPoly San Luis Obispo
netprl3
ION
Internet

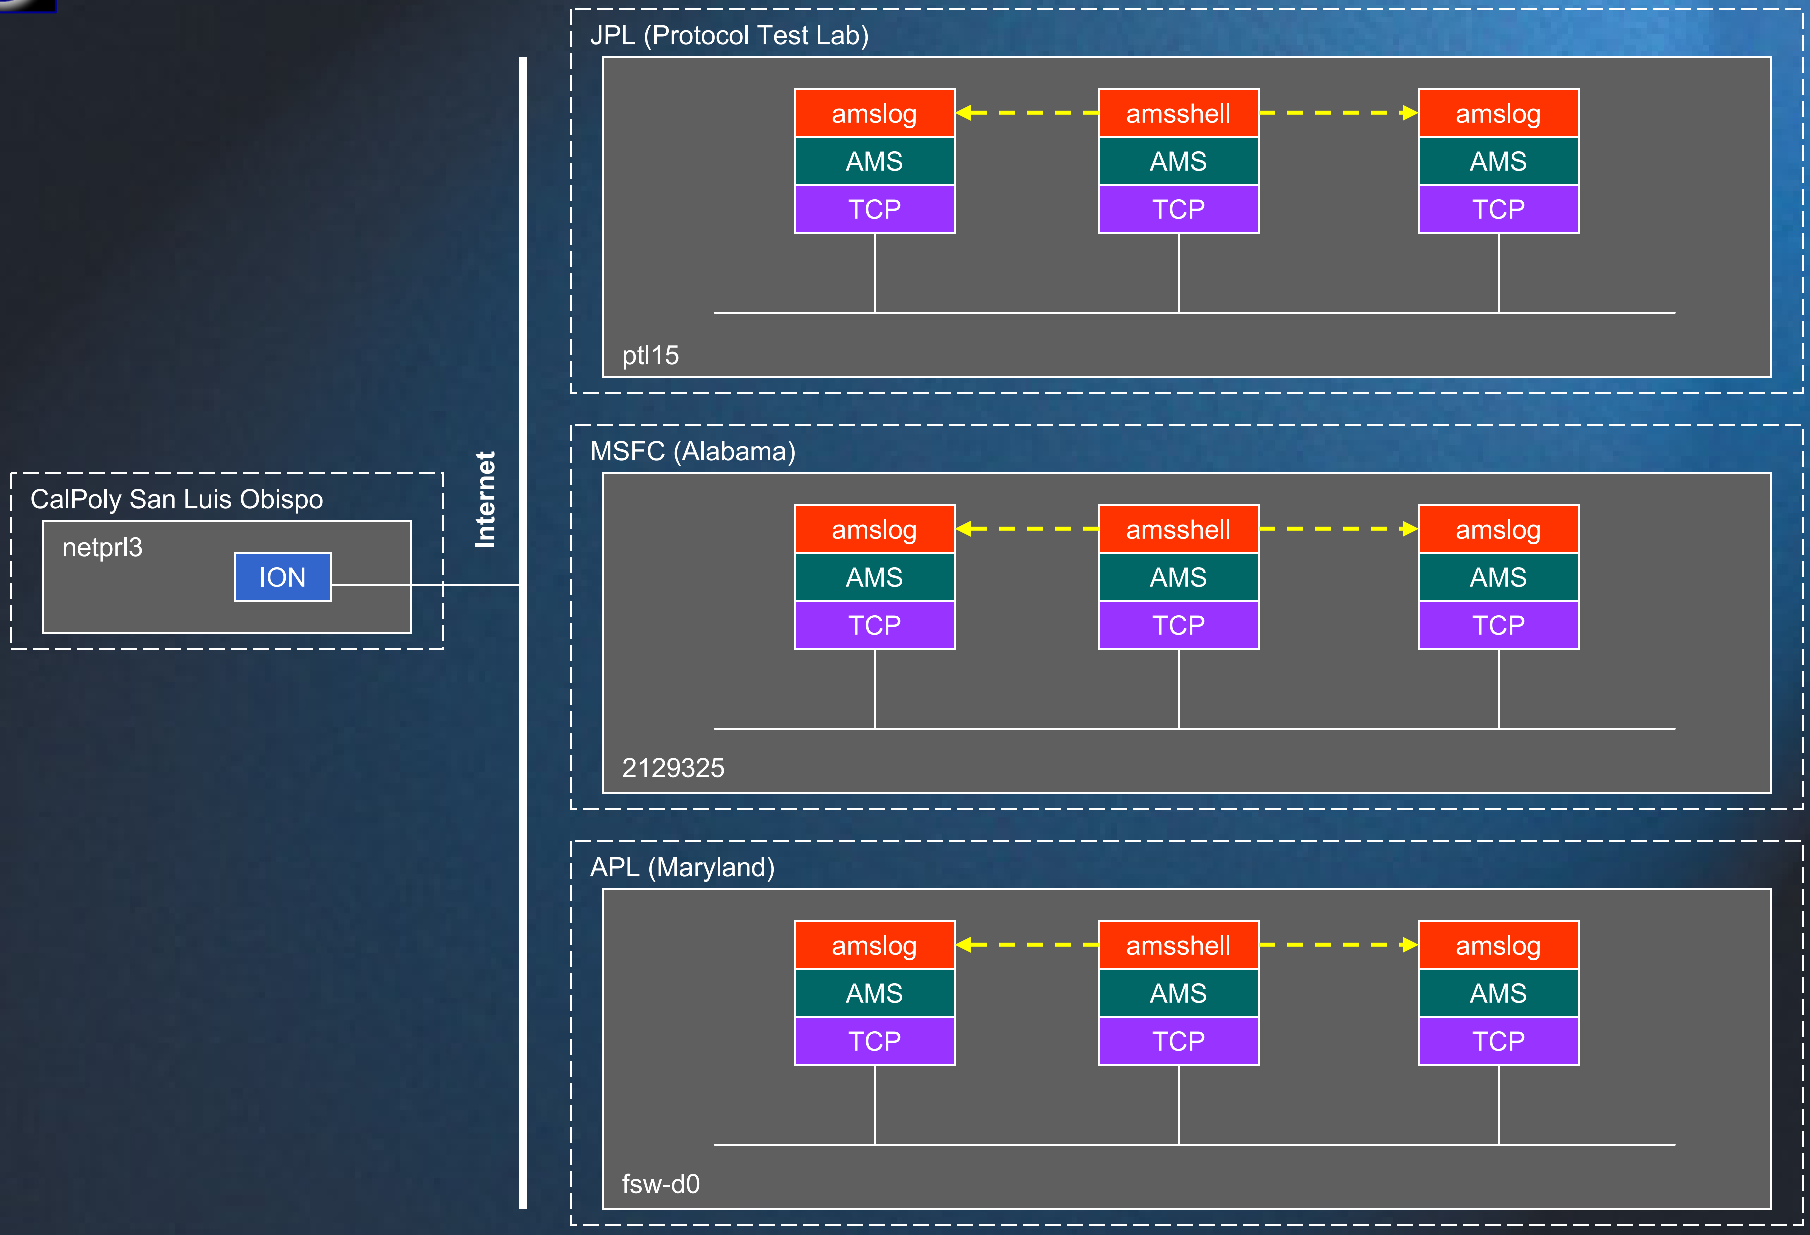

CalPoly San Luis Obispo
netprl3
ION
Internet
JPL (Protocol Test Lab)
amslog
AMS
TCP
amsshell
AMS
TCP
amslog
AMS
TCP
ptl15
MSFC (Alabama)
amslog
AMS
TCP
amsshell
AMS
TCP
amslog
AMS
TCP
2129325
APL (Maryland)
amslog
AMS
TCP
amsshell
AMS
TCP
amslog
AMS
TCP
fsw-d0

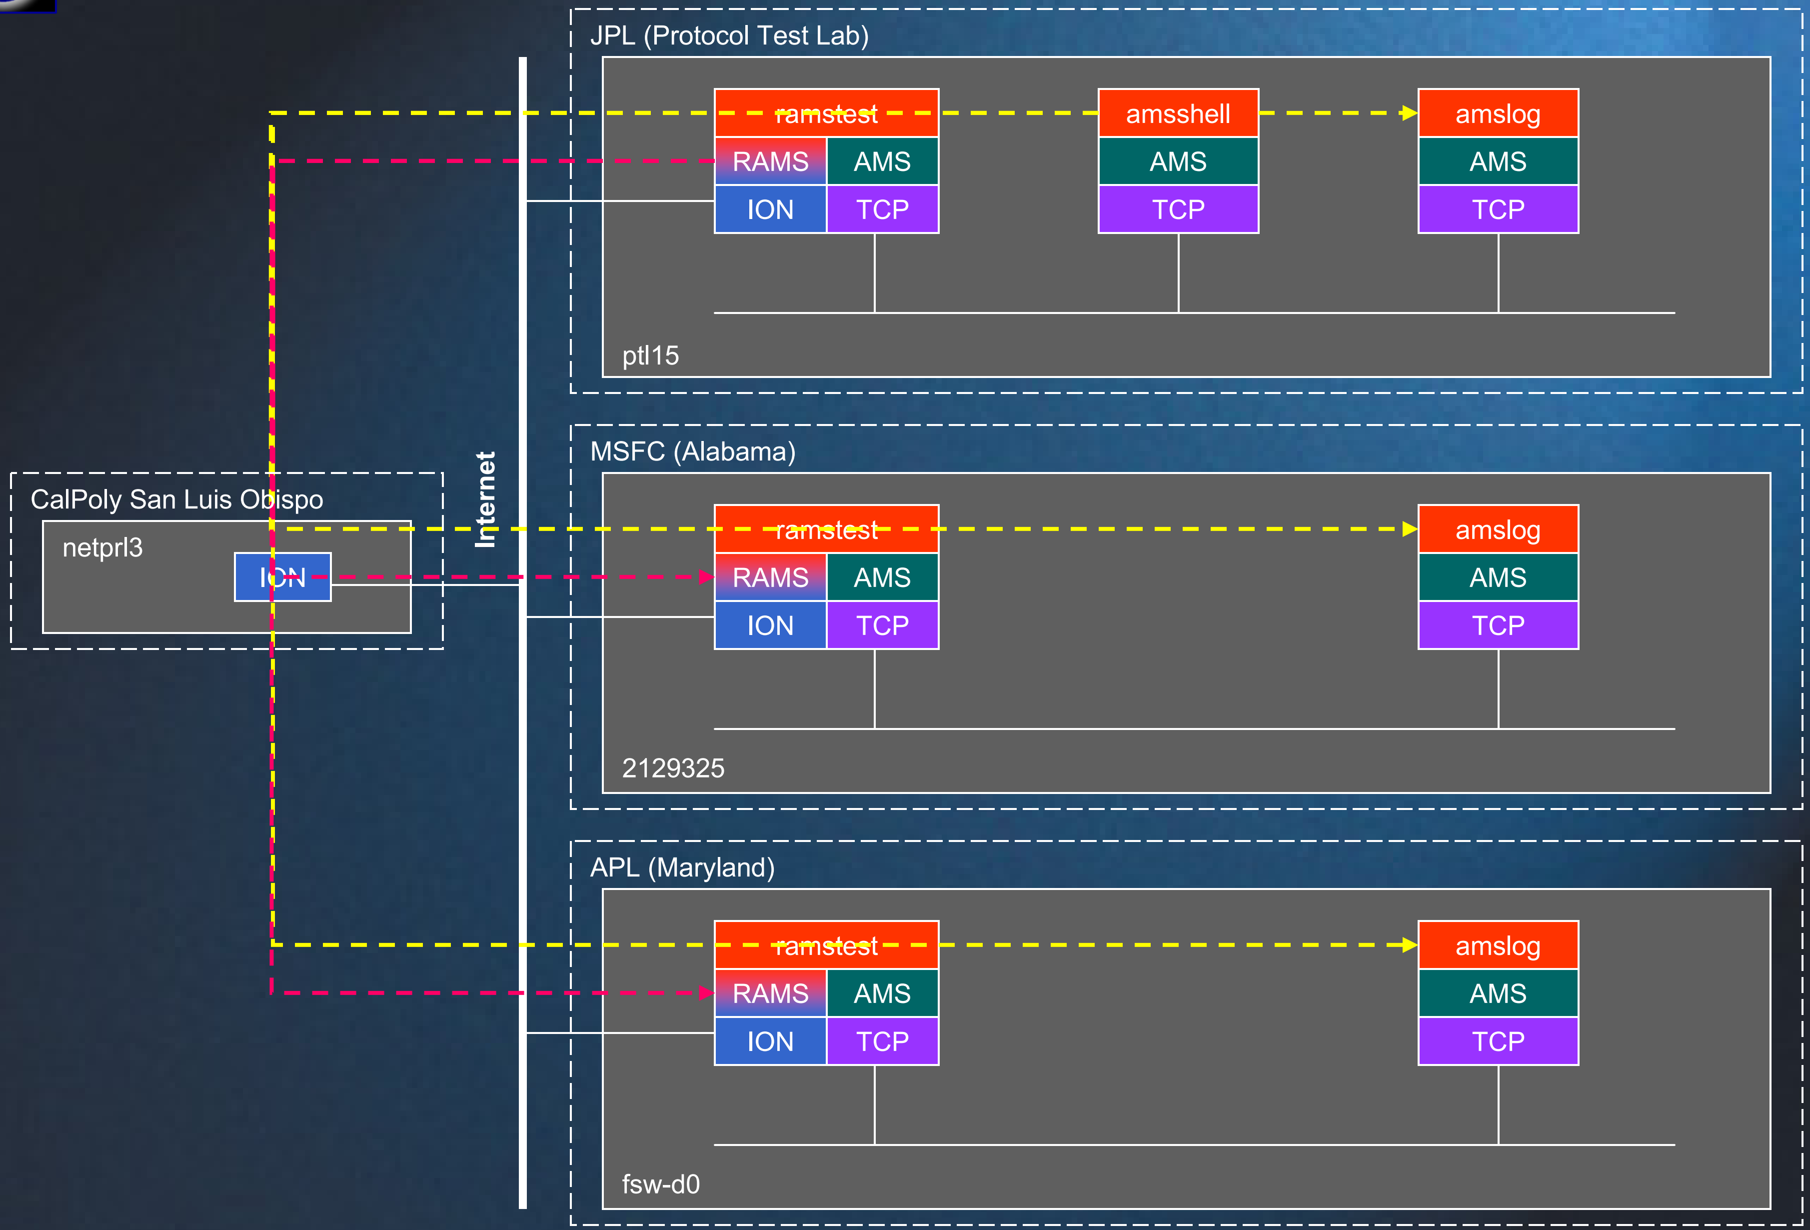

JPL (Protocol Test Lab)
ramstest
RAMS
AMS
ION
TCP
amsshell
AMS
TCP
amslog
AMS
TCP
ptl15
MSFC (Alabama)
ramstest
RAMS
AMS
ION
TCP
amslog
AMS
TCP
2129325
APL (Maryland)
ramstest
RAMS
AMS
ION
TCP
amslog
AMS
TCP
fsw-d0
CalPoly San Luis Obispo
netprl3
ION
Internet

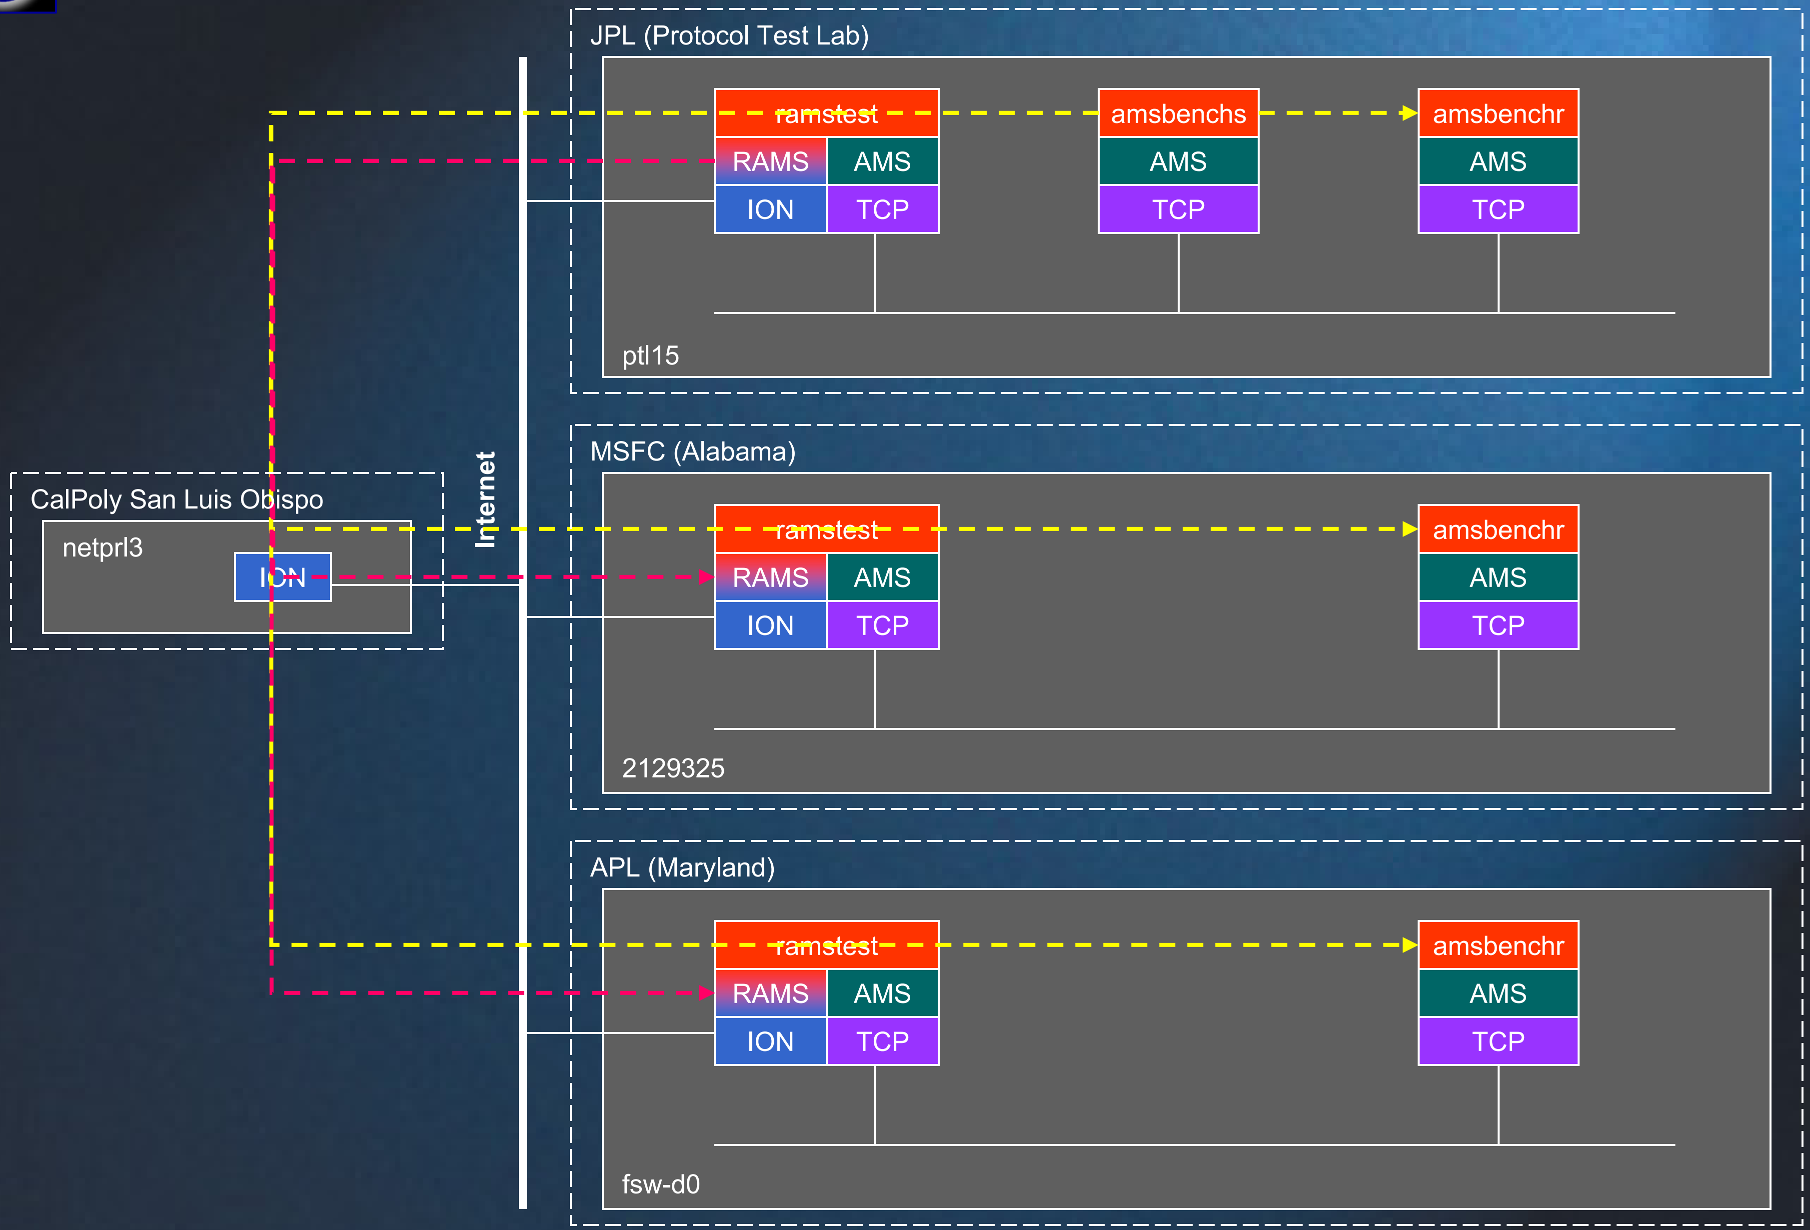

JPL (Protocol Test Lab)
ramstest
RAMS
AMS
ION
TCP
amsbenchs
AMS
TCP
amsbenchr
AMS
TCP
ptl15
CalPoly San Luis Obispo
netprl3
ION
Internet
MSFC (Alabama)
ramstest
RAMS
AMS
ION
TCP
amsbenchr
AMS
TCP
2129325
APL (Maryland)
ramstest
RAMS
AMS
ION
TCP
amsbenchr
AMS
TCP
fsw-d0

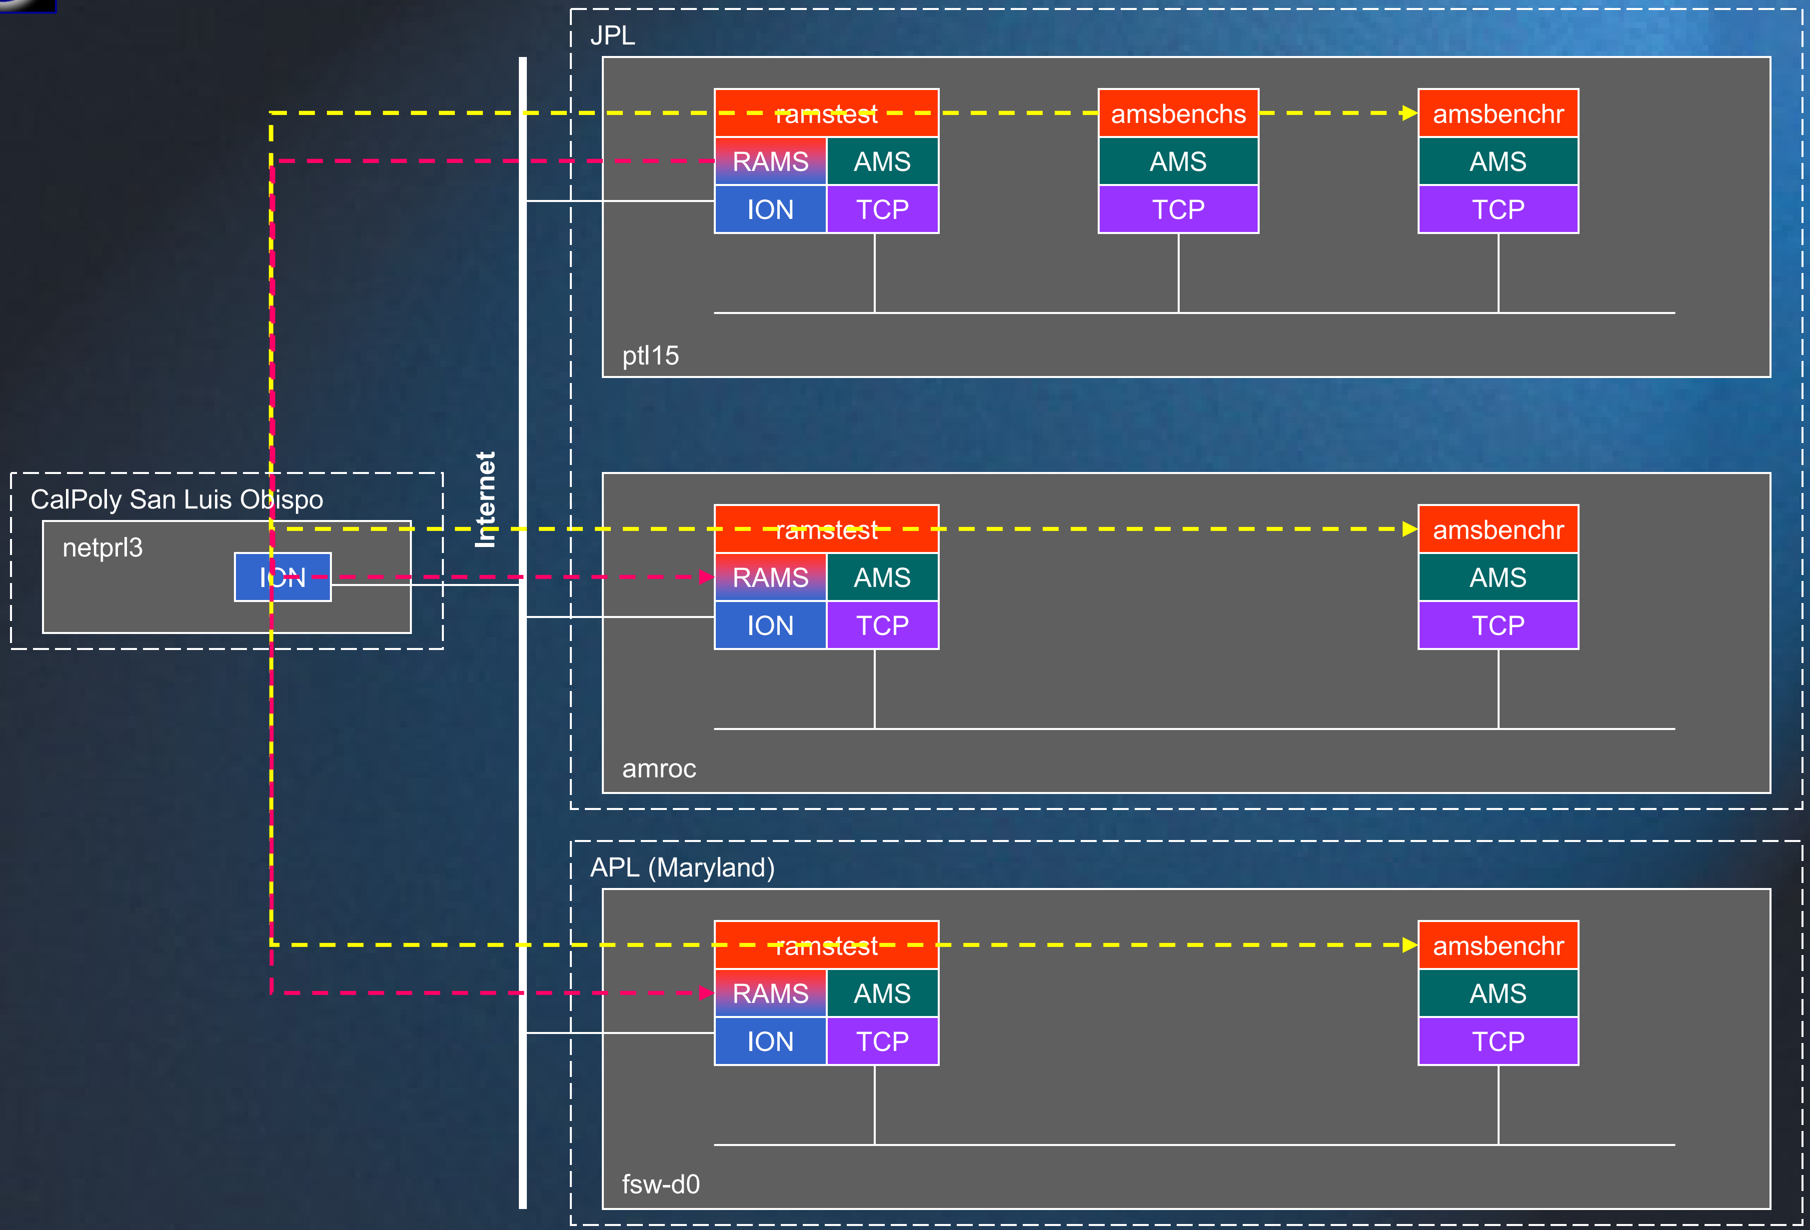

JPL
ramstest
RAMS
AMS
ION
TCP
amsbenchs
AMS
TCP
amsbenchr
AMS
TCP
ptl15
ramstest
RAMS
AMS
ION
TCP
amsbenchr
AMS
TCP
amroc
CalPoly San Luis Obispo
netprl3
ION
Internet
APL (Maryland)
ramstest
RAMS
AMS
ION
TCP
amsbenchr
AMS
TCP
fsw-d0

Results

- Successful exchange of AMS messages among all three centers.
- But couldn't maintain ION connectivity to MSFC for more than a few minutes.
  - Firewall measures at Marshall?
  - Still investigating.
- Performance benchmarking, with amsbenchs at one JPL machine:
  - Results at second JPL machine, simulating MSFC:
    - `Received 100 messages, totaling 2000000 bytes, in 2.037391 seconds.`
      `49.082 messages per second.`
      `7.489 Mbps.`
  - Results at APL:
    - `Received 100 messages, totaling 2000000 bytes, in 2.056565 seconds.`
      `48.625 messages per second.`
      `7.420 Mbps.`

Current Status

- Protocol specifications seem mostly mature. “Draft Standard” issued for review by CCSDS Area Manager.
- Only one implementation so far. Second implementation needed, to drive out problems in specs.
- Most features of the protocols have been implemented, but much testing remains to be done.
- Currently testing port to VxWorks.
  - Intent is to support real-time message exchange in embedded systems.
- RAMS needs performance optimization.
- Documentation still very primitive.
- Reference implementation has been distributed to NASA centers (Goddard, Glenn, Ames, Marshall, Johnson), APL, ESA, CNES, MITRE, Ohio University, APL, Cal Poly San Luis Obispo.