IP SHOWCASE™ THEATER

CURATED BY VSF VIDEO SERVICES FORUM

Scalability and Performance of IS-04 and IS-05 and How TR-1001-1 Helps

Rob Porter

Sony Europe B.V.

IP SHOWCASE™ THEATER

AMWA NMOS IS-04 and IS-05 APIs

The AMWA NMOS Scalability Study

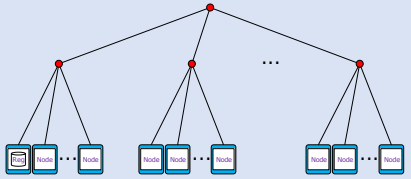

Scalability Study Results

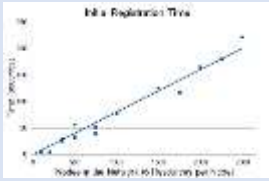

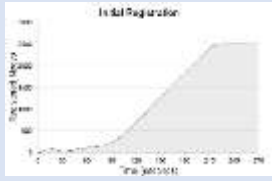

Best Practice Recommendations
and How JT-NM TR-1001-1 Helps

SONY

IP SHOWCASE™ THEATER

AMWA NMOS IS-04 and IS-05 APIs

The AMWA NMOS Scalability Study

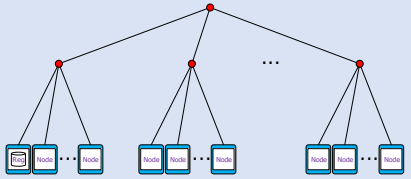

Scalability Study Results

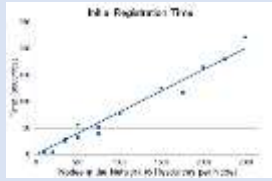

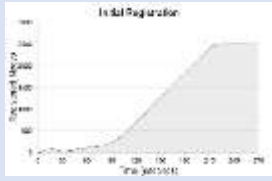

Best Practice Recommendations
and How JT-NM TR-1001-1 Helps

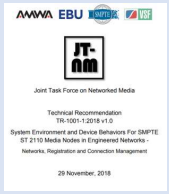

SONY

IP SHOWCASE™ THEATER

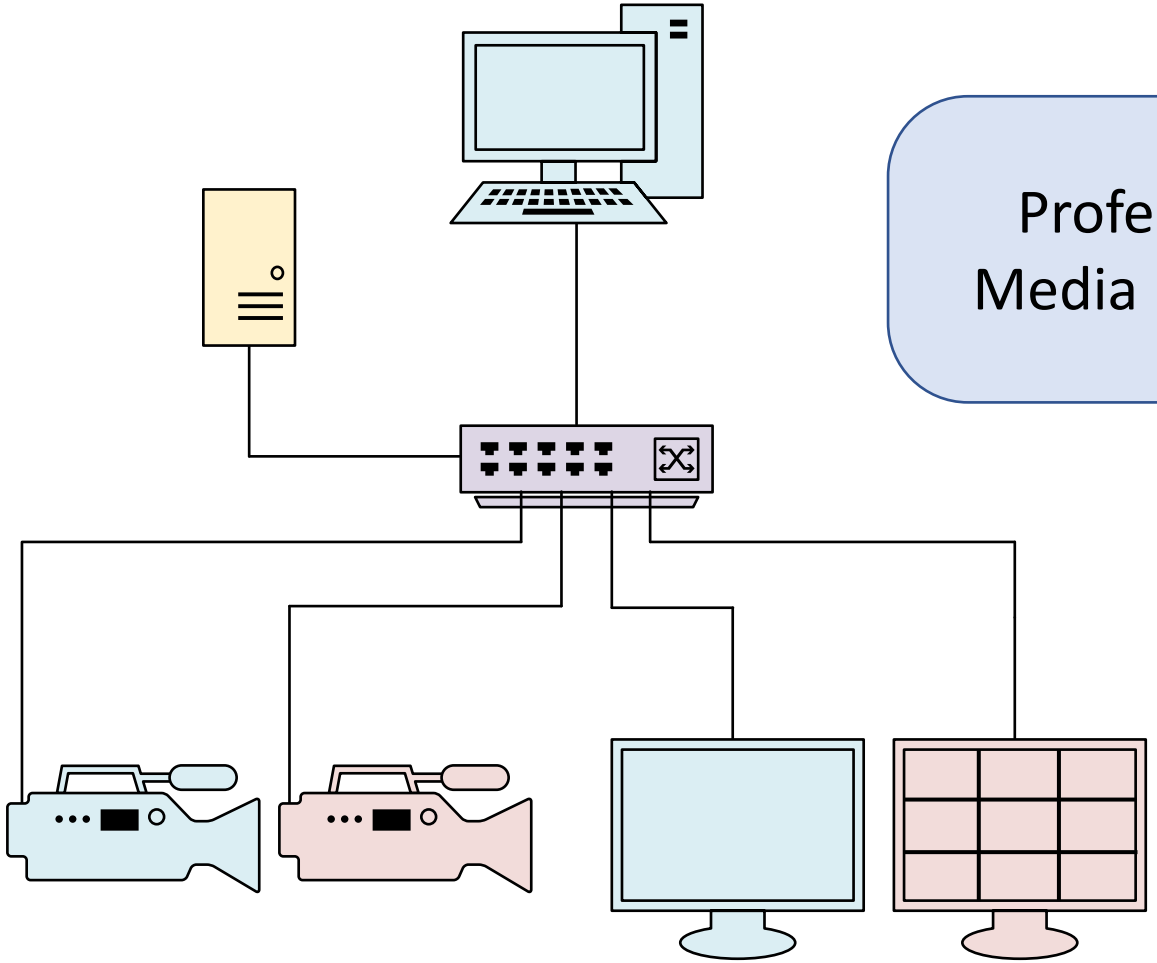

Professional
Media Network

SONY

IP SHOWCASE™ THEATER

Network Infrastructure

- One or more network switches
- Any architecture
  - Single monolithic switch
  - Multi-layer spine-leaf network

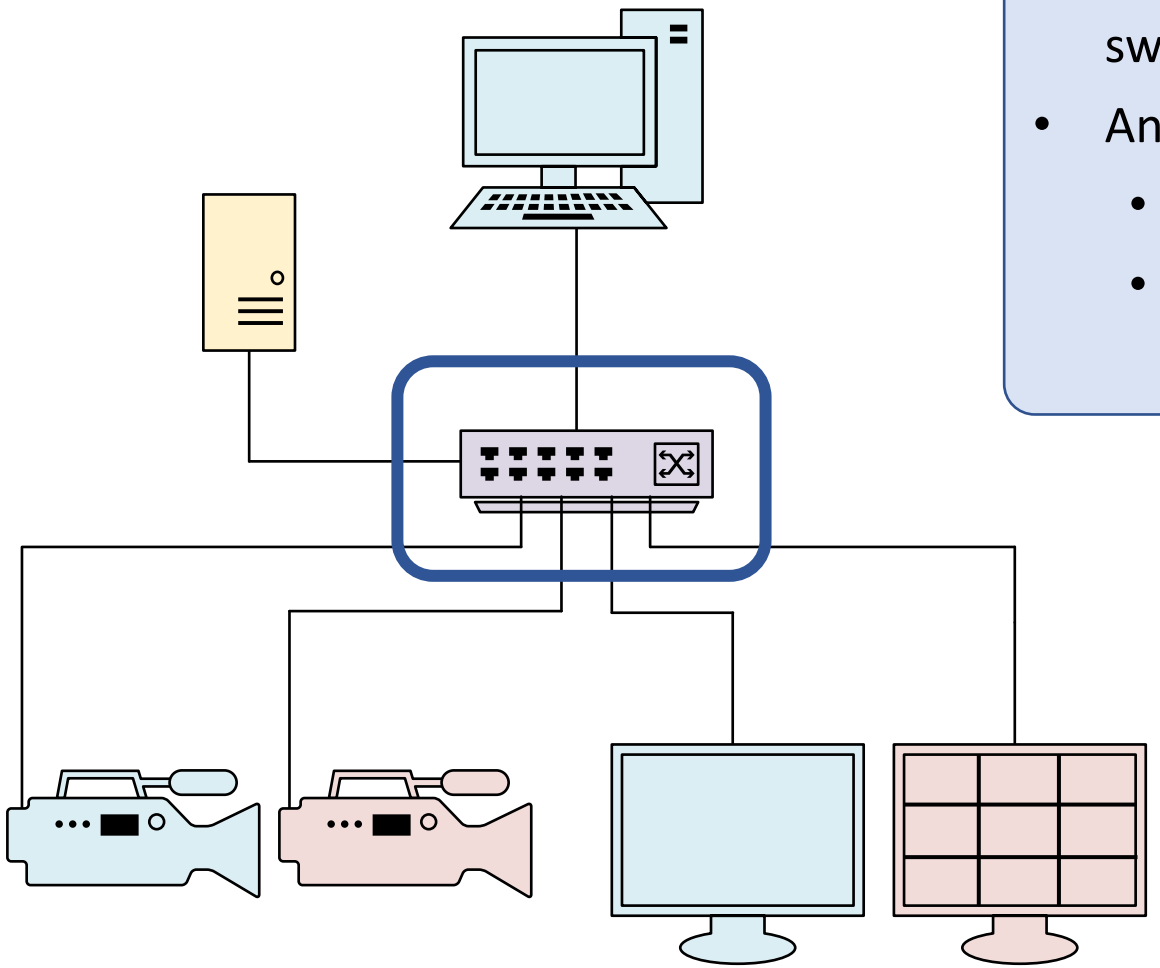

SONY

IP SHOWCASE™ THEATER

Media Nodes

- **Node** is a logical host connected to the network
- Can host one or more **Devices** each with any number of associated **Senders**, **Receivers**, **Sources** and **Flows**
- These are known as *resources* and are all defined by the JT-NM content model

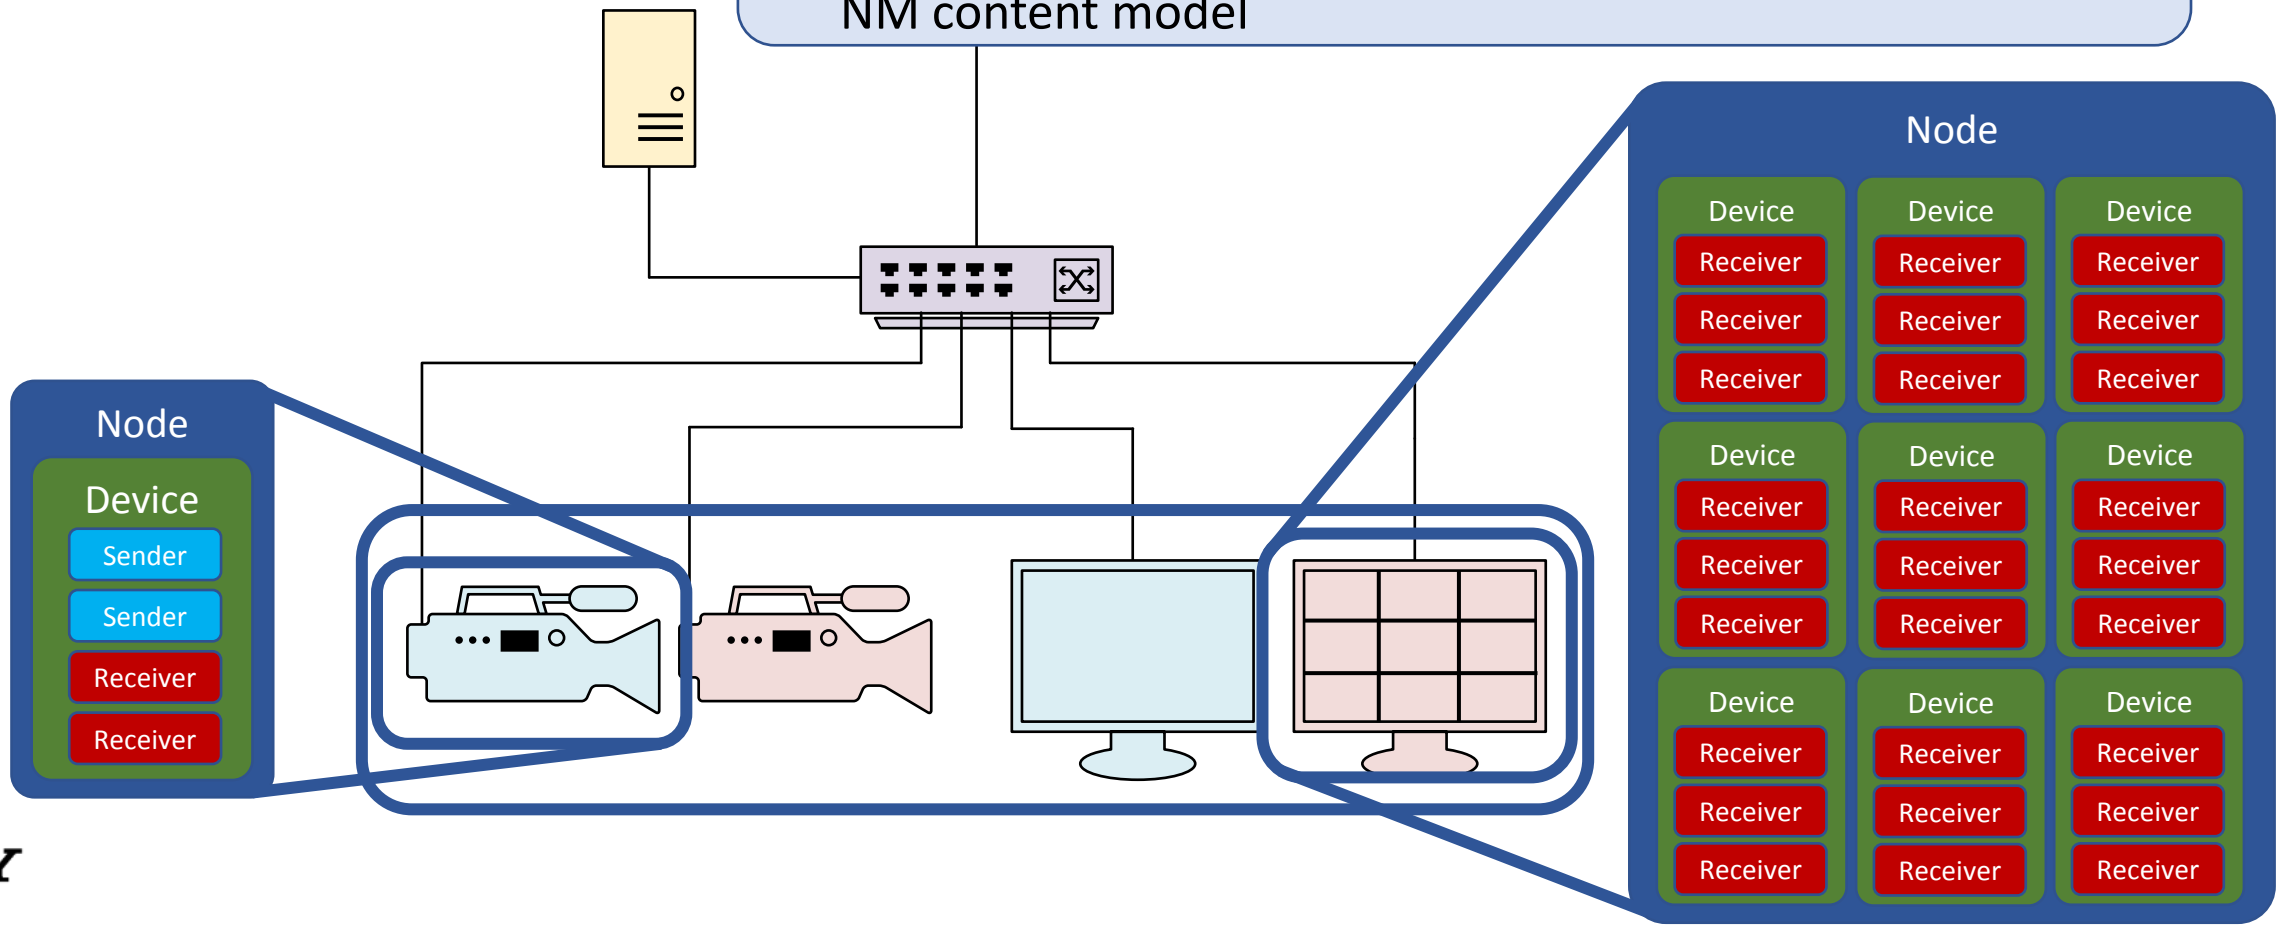

SONY

IP SHOWCASE™ THEATER

IS-04 Registration and Discovery System (RDS)

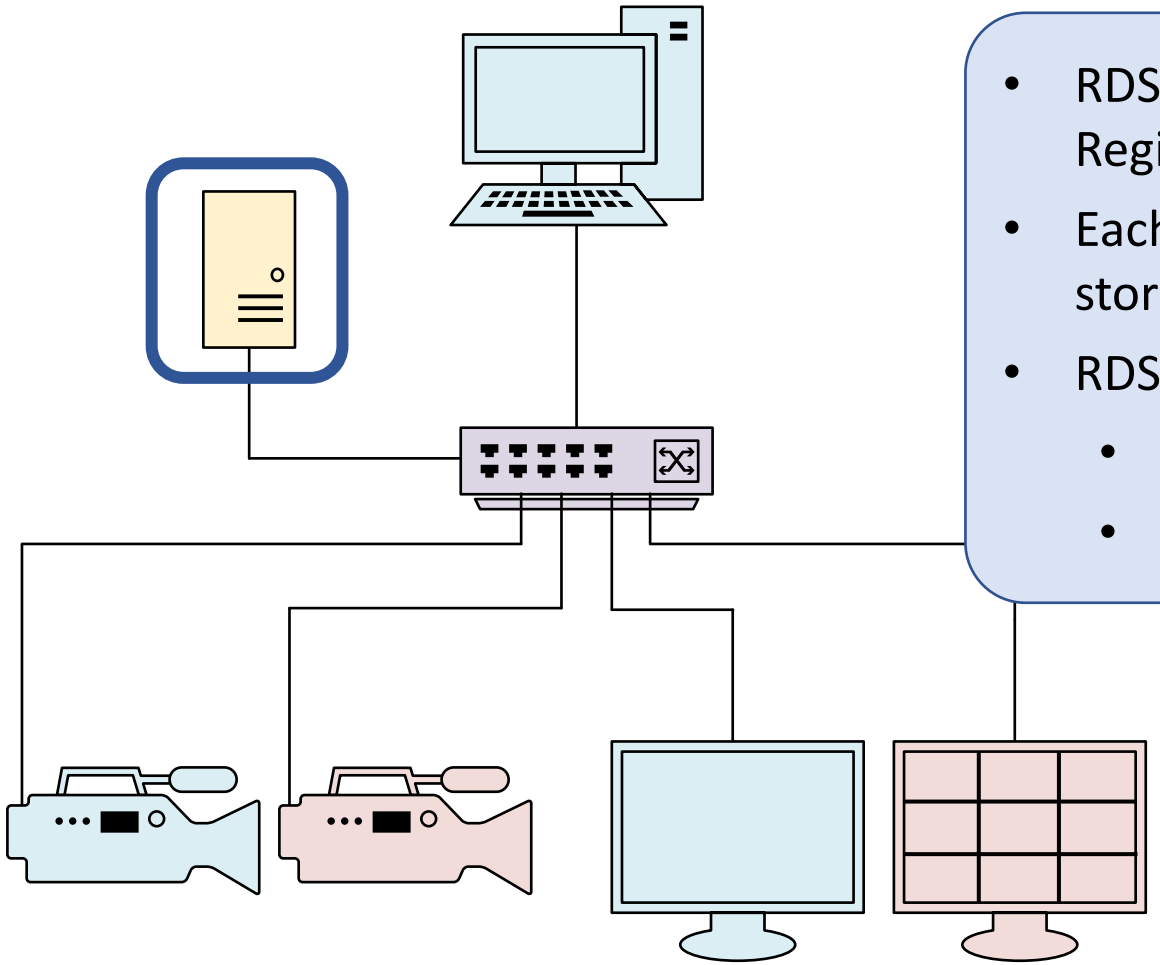

- RDS comprises one or more Registry instances
- Each includes a database storing all registered resources
- RDS exposes two APIs:
  - Registration API
  - Query API

SONY

IP SHOWCASE™ THEATER

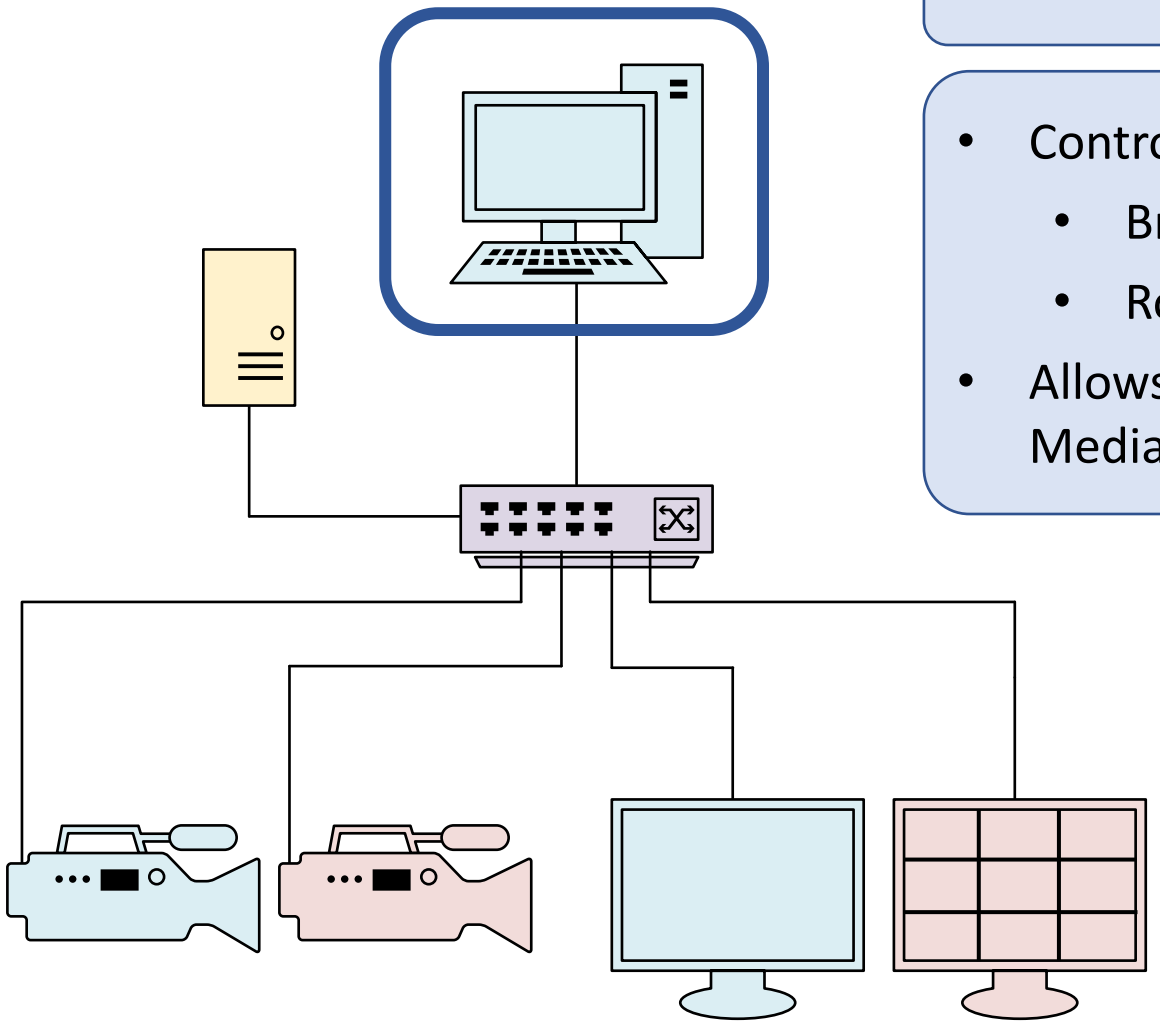

NMOS Client

- Control system, e.g.
  - Broadcast Controller
  - Routing panel
- Allows connections between Media Nodes to be set up

SONY

IS-04 Registration API

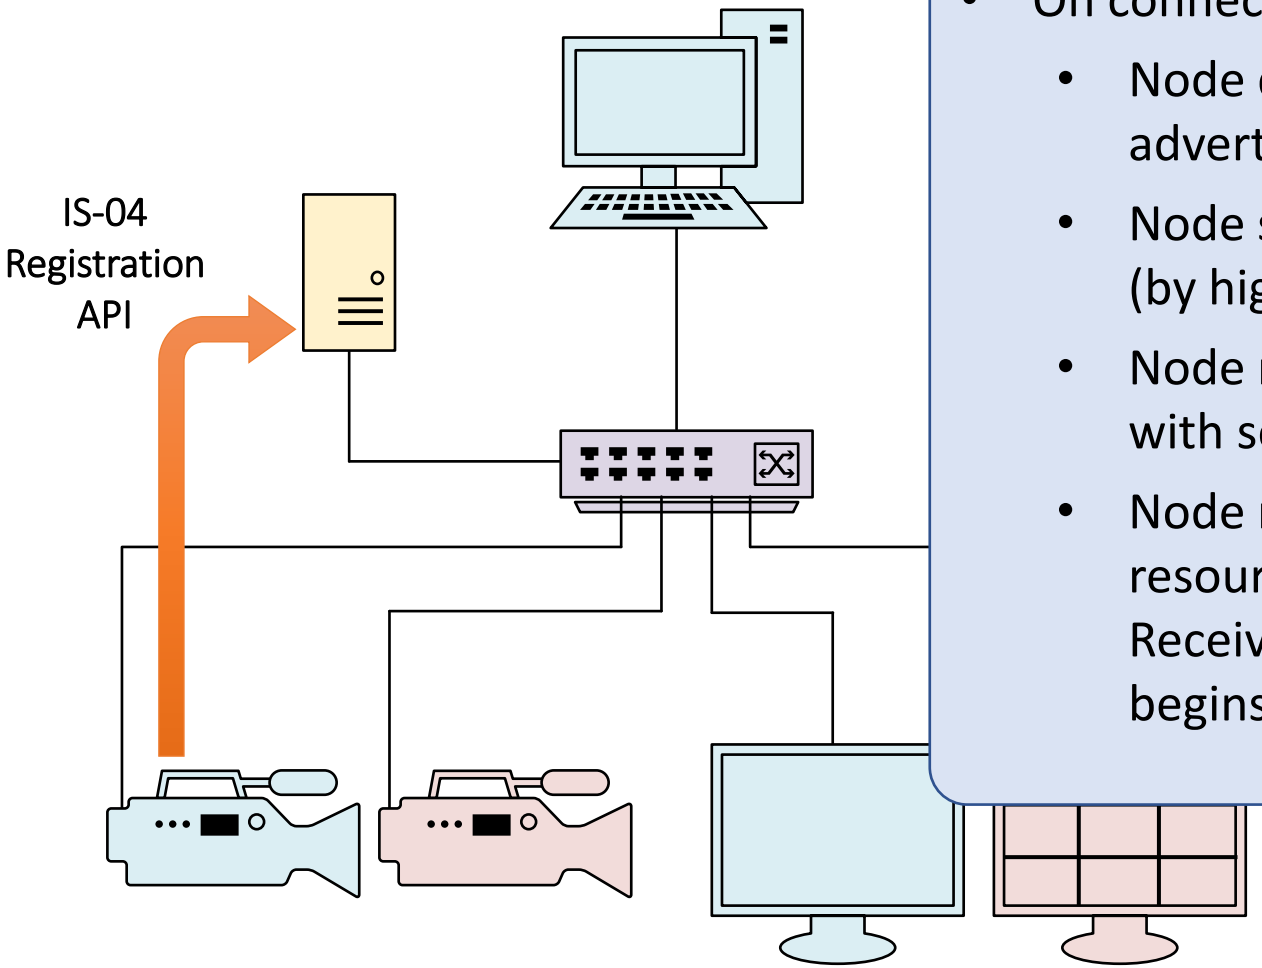

- On connecting a Node to the network:
  - Node discovers Registration APIs advertised over DNS-SD
  - Node selects a Registration API (by highest priority)
  - Node registers its Node resource with selected Registration API
  - Node registers each of its sub-resources (Devices, Senders, Receivers, Sources, Flows) and begins to post regular heartbeats

Registration API Request (Node→RDS):

```
POST /x-nmos/registration/v1.2/resource HTTP/1.1
Content-Type: application/json
 {
   "type": "node",
   "data": {
     "version": "1441973902:879053935",
     "hostname": "host1",
     "label": "host1",
     "description": "host1",
     "tags": {},
     "href": "http://172.29.80.65:12345/",
     "api": {
       "versions": ["v1.1", "v1.2"],
       "endpoints": [
         {
           "host": "172.29.80.65",
           "port": 12345,
           "protocol": "http"
         },
         {
           "host": "172.29.80.65",
           "port": 443,
           "protocol": "https"
         }
       ]
     },
```

Node

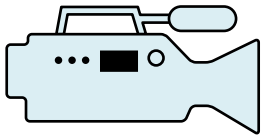

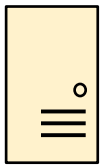

Registry

SONY

IP SHOWCASE™ THEATER

Registration API Request (Node→RDS):

```
POST /x-nmos/registration/v1.2/resource HTTP/1.1
Content-Type: application/json
 {
   "type": "node",
   "data": {
     "version": "1441973902:879053935",
     "hostname": "host1",
     "label": "host1",
     "description": "host1",
     "tags": {},
     "href": "http://172.29.80.65:12345/",
     "api": {
       "versions": ["v1.1", "v1.2"],
       "endpoints": [
         {
           "port": 443,
           "protocol": "https"
         }
       ]
     },
```

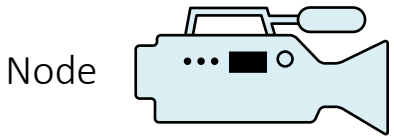

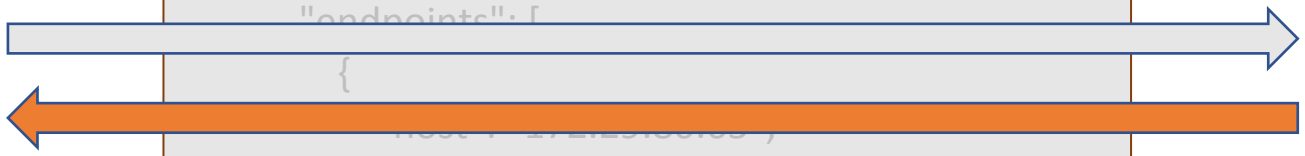

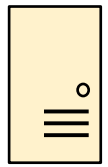

Registration API Response (RDS→Node):
HTTP/1.1 201 Created
/x-nmos/registration/v1.2/resource/nodes/3b8be755-08ff-452b-b217-c9151eb21193/

SONY

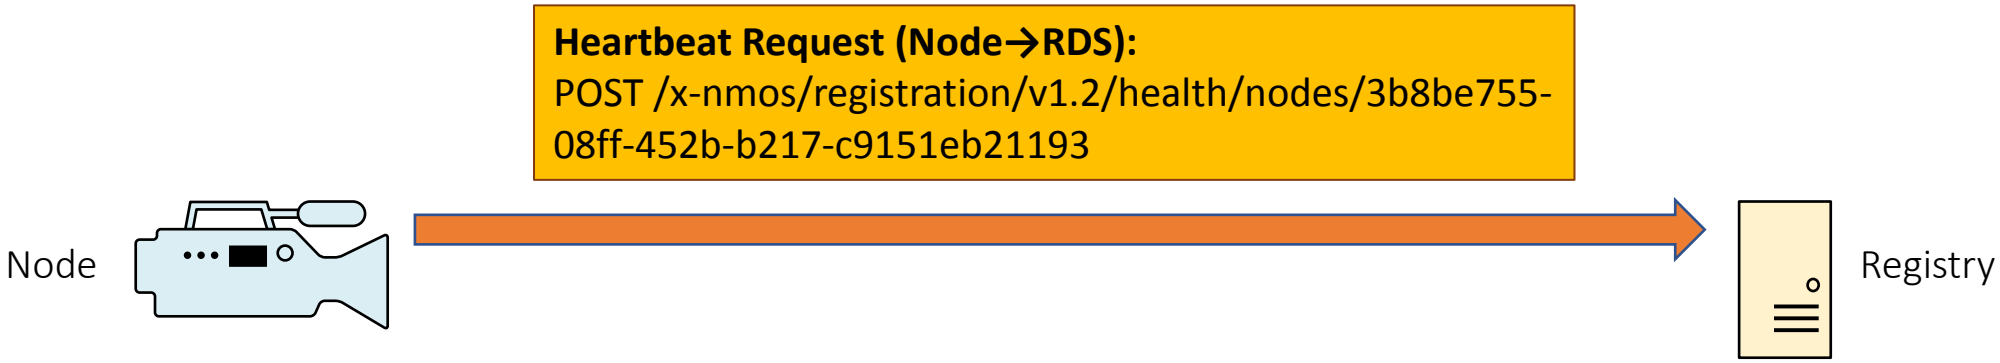
Heartbeat Request (Node→RDS):
POST /x-nmos/registration/v1.2/health/nodes/3b8be755-08ff-452b-b217-c9151eb21193
Node
Registry

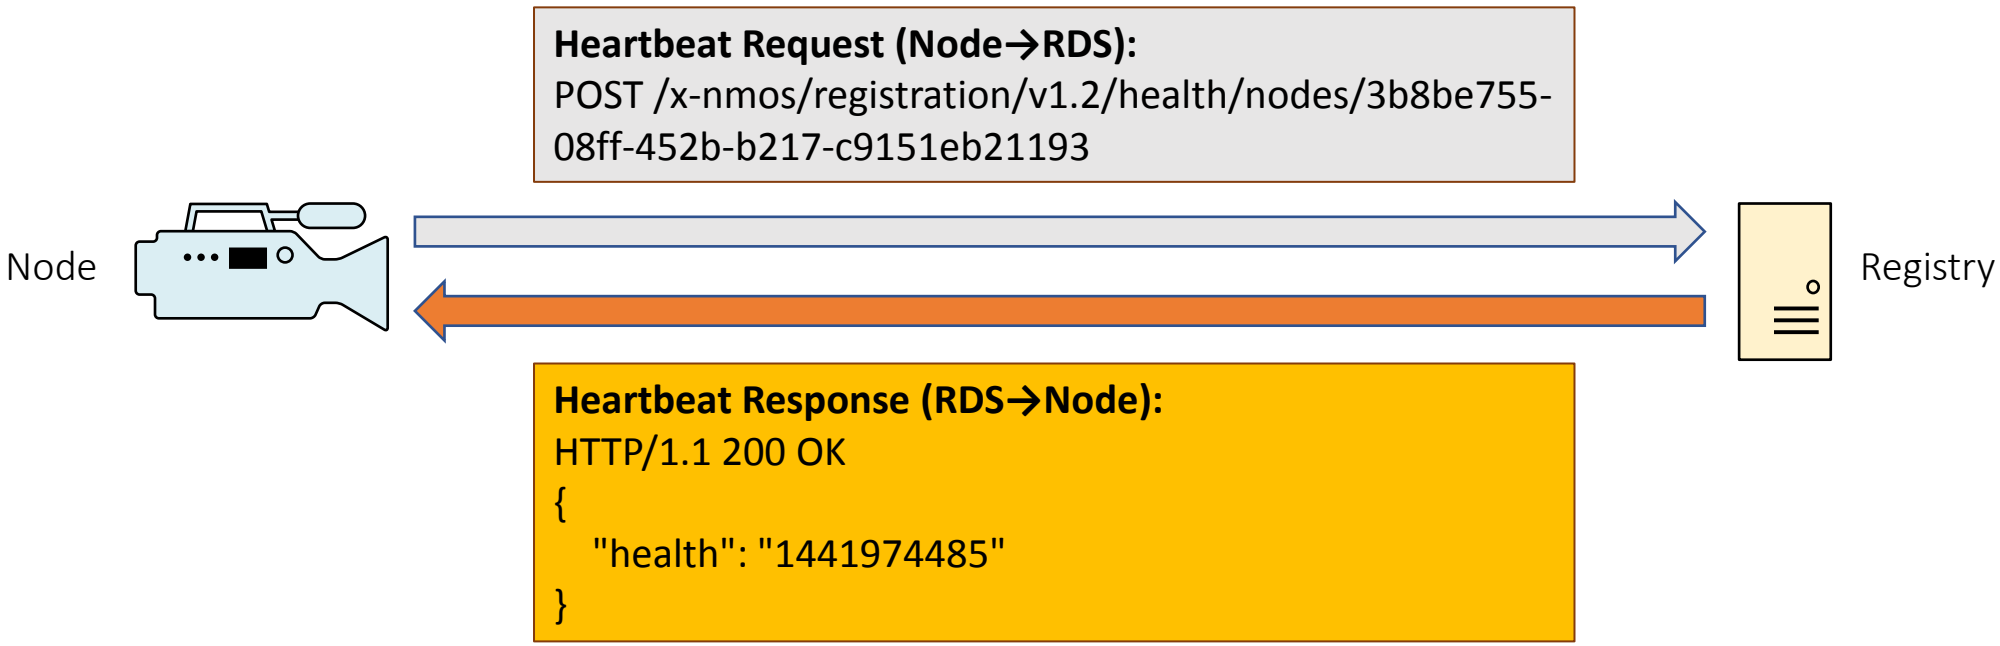
Heartbeat Request (Node→RDS):
POST /x-nmos/registration/v1.2/health/nodes/3b8be755-08ff-452b-b217-c9151eb21193
Node
Registry
Heartbeat Response (RDS→Node):
HTTP/1.1 200 OK
{
"health": "1441974485"
}

IP SHOWCASE™ THEATER

IS-04 Query API

- Client is able to get list of registered resources from Registry using IS-04 Query API
- It can also subscribe to WebSocket notifications of changes in the RDS

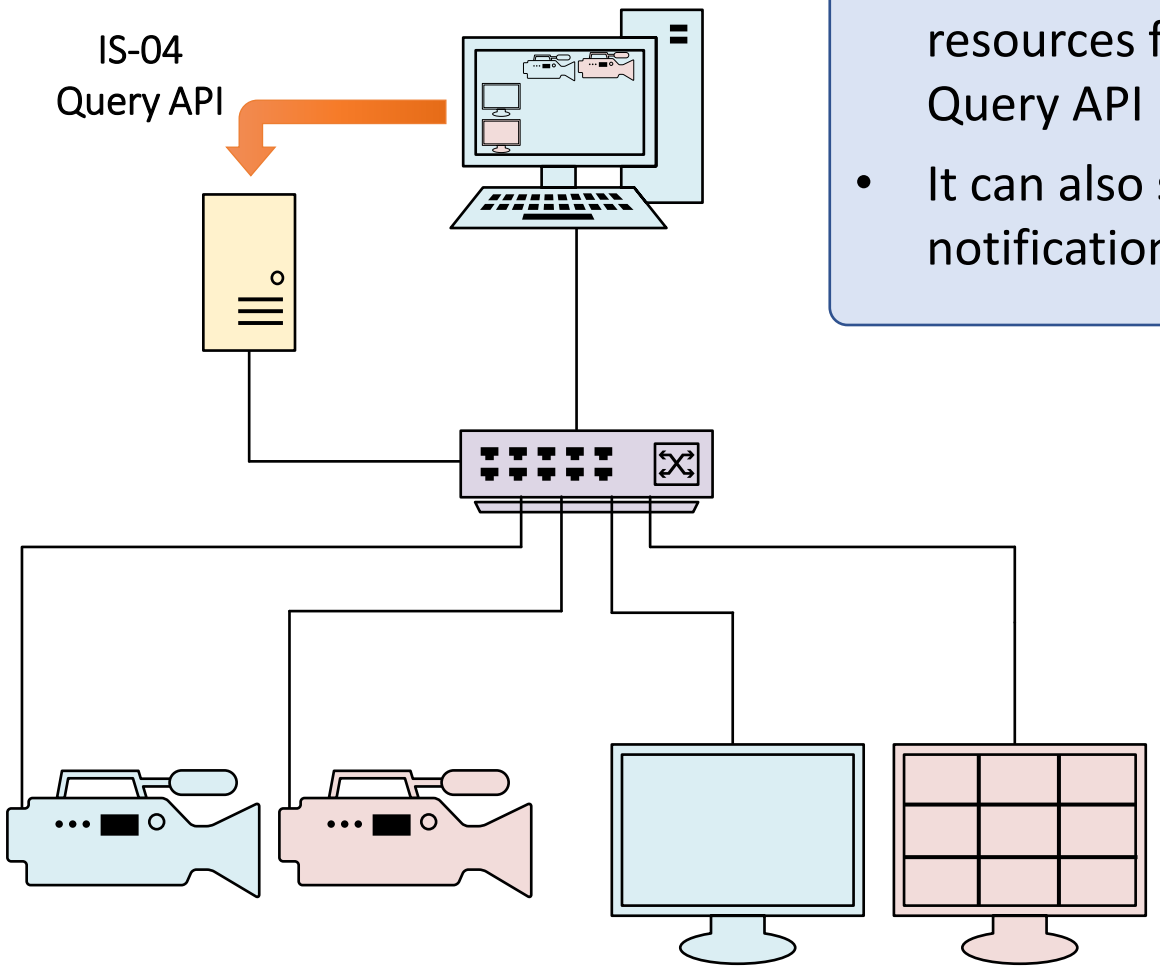

SONY

IP SHOWCASE™ THEATER

IS-05 Connection API

- Client can use a Node's IS-05 Connection API to make a connection
  - Get transport file from Sender
  - Patch transport file to Receiver
  - Join multicast group

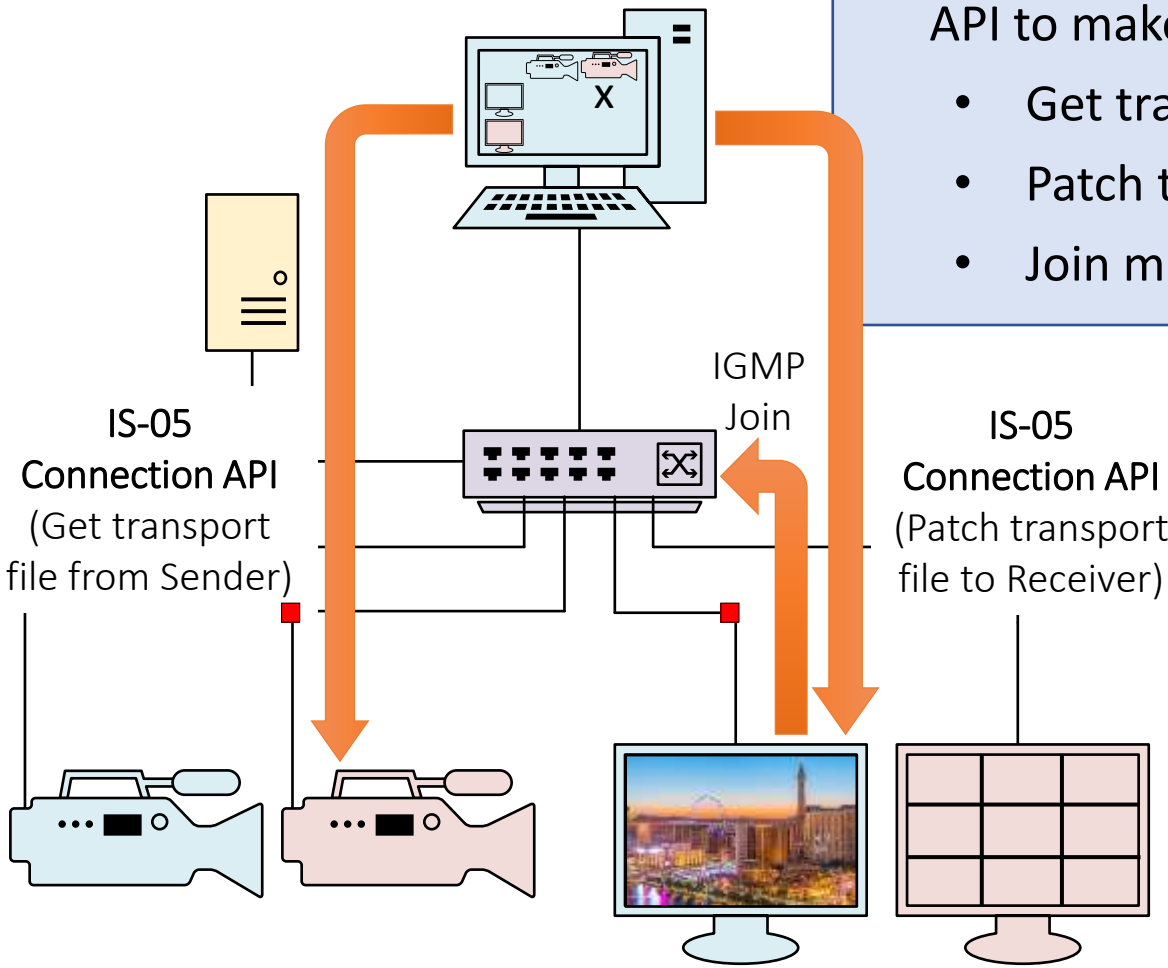

SONY

IP SHOWCASE™ THEATER

Transport file (SDP file):

v=0
o=- 3755583281 3755583281 IN IP4 192.168.9.142
s=Camera 1 Video
t=0 0
a=group:DUP PRIMARY SECONDARY
m=video 50020 RTP/AVP 96
c=IN IP4 239.22.142.1/32
a=ts-refclk:ptp=IEEE1588-2008:traceable
a=mediaclk:direct=0
a=source-filter: incl IN IP4 239.22.142.1 192.168.9.142
a=rtpmap:96 raw/90000
a=fmtp:96 width=1920; height=1080; exactframerate=30000/1001;
interlace; sampling=YCbCr-4:2:2; depth=10; colorimetry=BT709;
TCS=SDR; PM=2110GPM; SSN=ST2110-20:2017; TP=2110TPN;
a=mid:PRIMARY
m=video 50120 RTP/AVP 96
c=IN IP4 239.122.142.1/32
a=ts-refclk:ptp=IEEE1588-2008:traceable
a=mediaclk:direct=0
a=source-filter: incl IN IP4 239.122.142.1 192.168.109.142
a=rtpmap:96 raw/90000
a=fmtp:96 width=1920; height=1080; exactframerate=30000/1001;
interlace; sampling=YCbCr-4:2:2; depth=10; colorimetry=BT709;
TCS=SDR; PM=2110GPM; SSN=ST2110-20:2017; TP=2110TPN;
a=mid:SECONDARY

Client

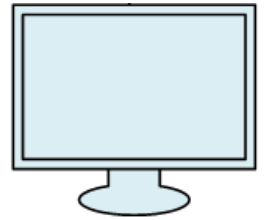

Node (Receiver)

SONY

IP SHOWCASE™ THEATER

AMWA NMOS IS-04 and IS-05 APIs

The AMWA NMOS Scalability Study

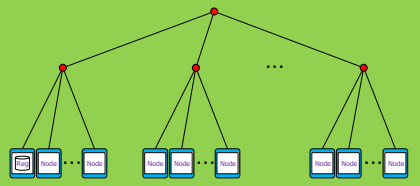

Scalability Study Results

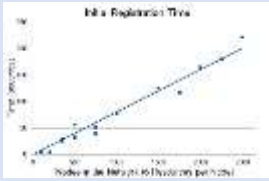

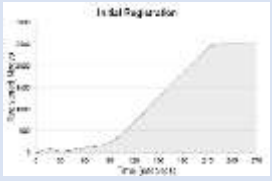

Best Practice Recommendations
and How JT-NM TR-1001-1 Helps

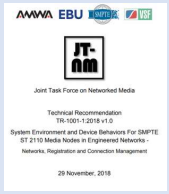

SONY

IP SHOWCASE™ THEATER

The AMWA NMOS Scalability Study

- A key requirement of the AMWA IS-04 and IS-05 APIs is that they can be used reliably at scale
  - i.e. for very large networks comprising thousands of NMOS Nodes such as might be found in a typical broadcast installation.
- The aim of the AMWA NMOS Scalability Study was to help address this
- Study took place within the AMWA community and was led by Sony
- The study used a virtualised network to test and make timing measurements of various IS-04 and IS-05 operations at scale

IP SHOWCASE™ THEATER

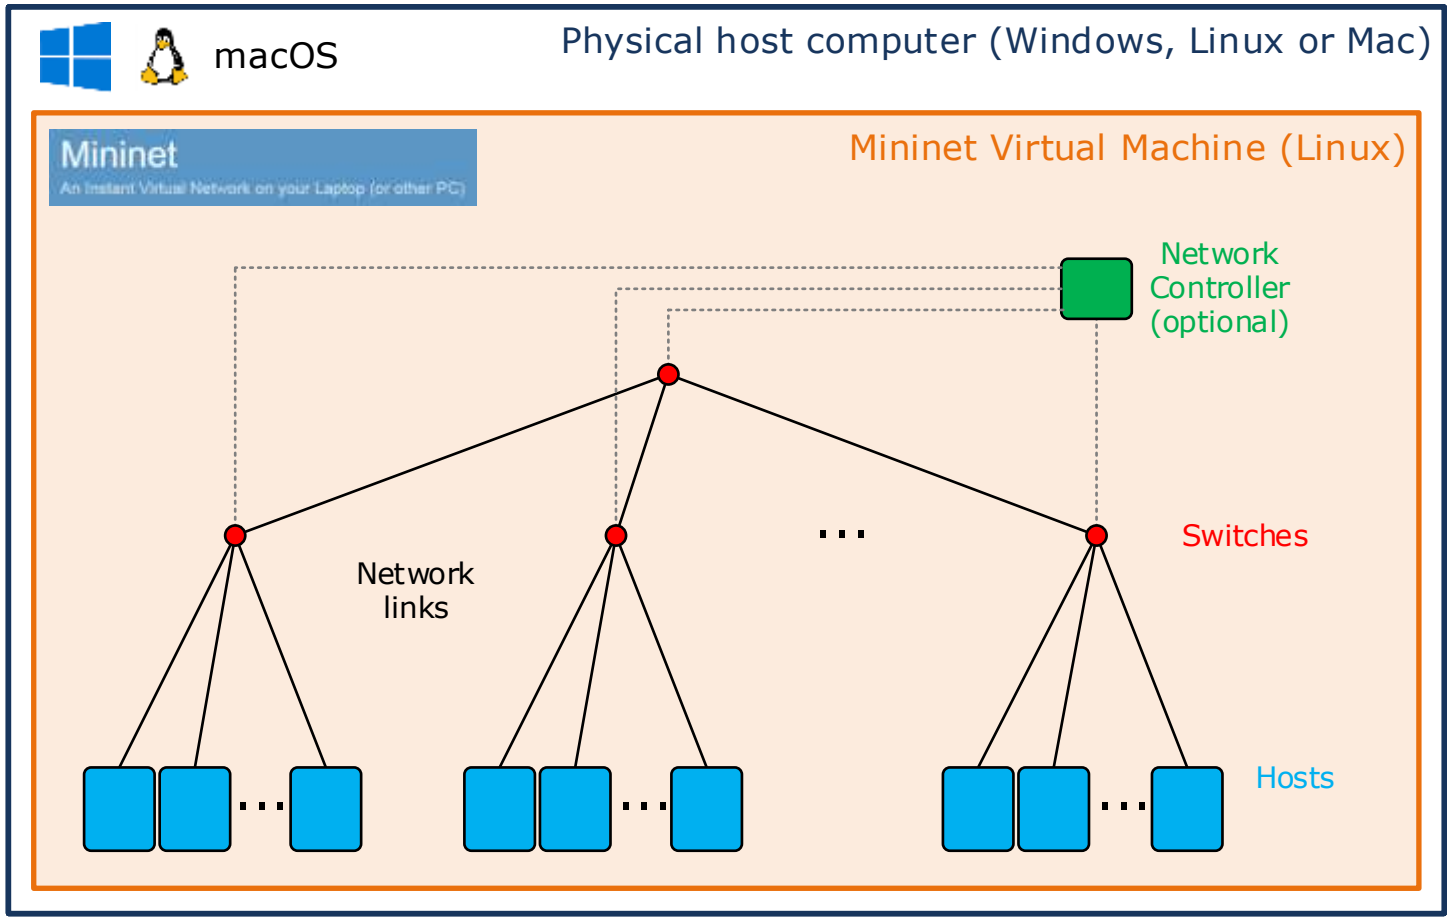

Scalability Study Methodology

- Use Mininet virtualised network to simulate large number of network endpoints
- Mininet extended for NMOS to allow NMOS processes to be run on each Mininet host

SONY

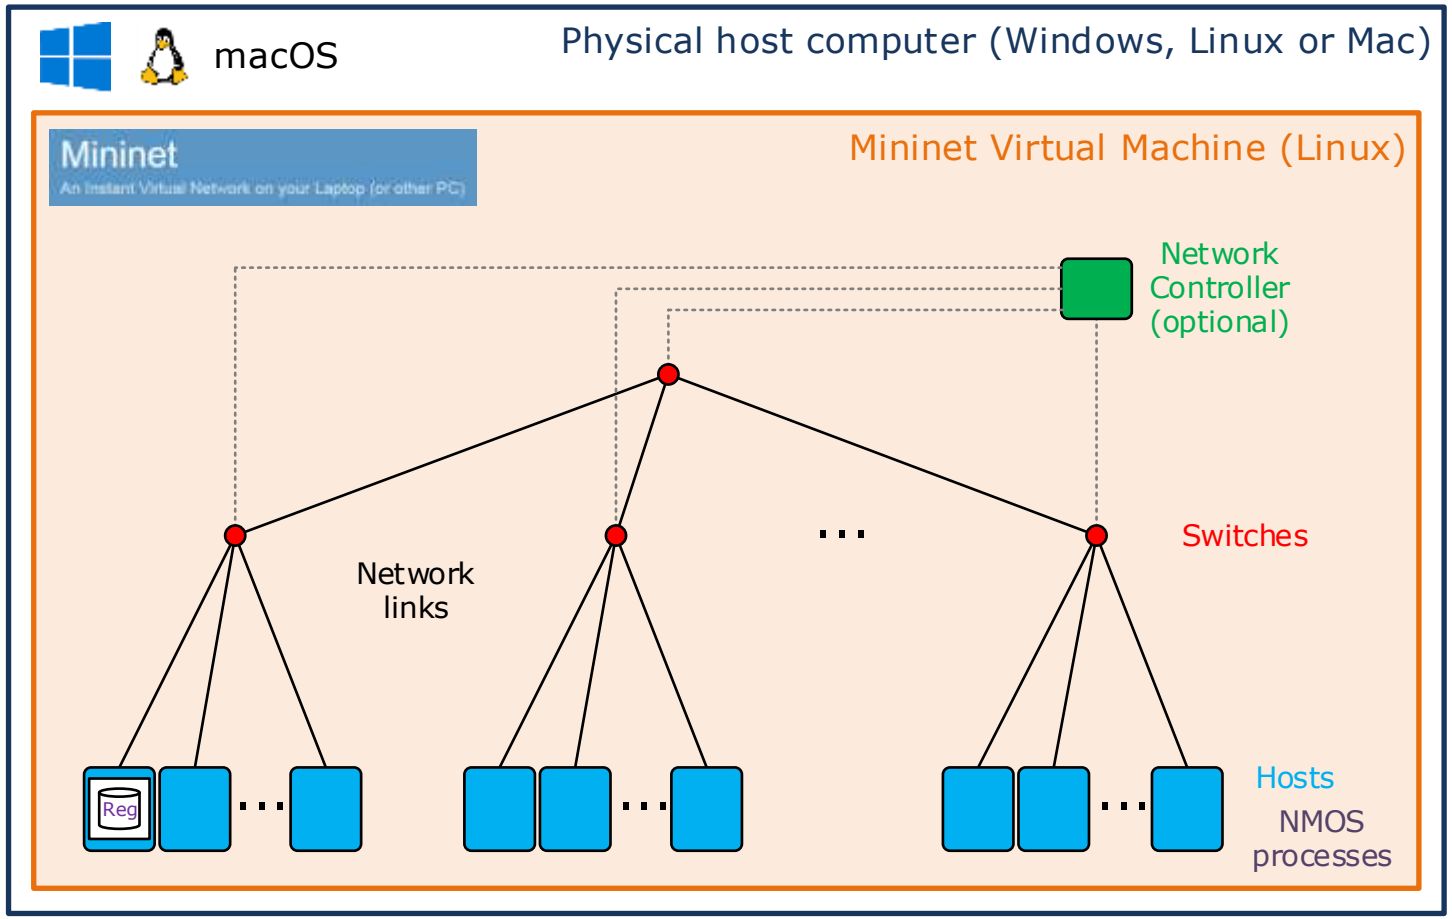

Scalability Study Methodology

- Use Mininet virtualised network to simulate large number of network endpoints
- Mininet extended for NMOS to allow NMOS processes to be run on each Mininet host
- Run nmos-cpp-registry* on one Mininet host

* https://github.com/sony/nmos-cpp

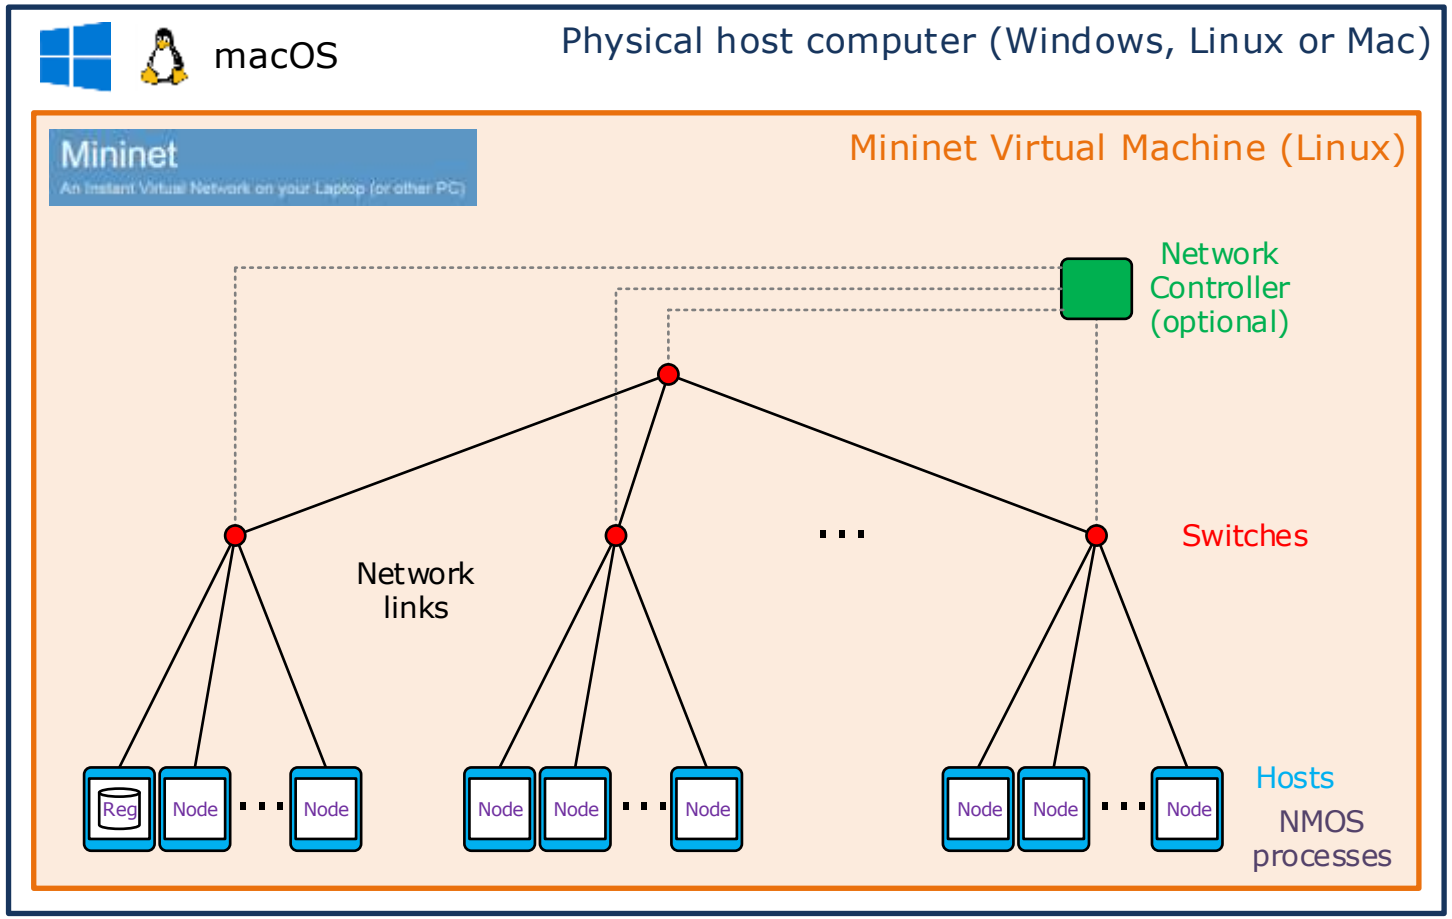

Scalability Study Methodology

- Use Mininet virtualised network to simulate large number of network endpoints
- Mininet extended for NMOS to allow NMOS processes to be run on each Mininet host
- Run nmos-cpp-registry* on one Mininet host
- Run multiple instances of nmos-cpp-node* on multiple other Mininet hosts

* https://github.com/sony/nmos-cpp

IP SHOWCASE™ THEATER

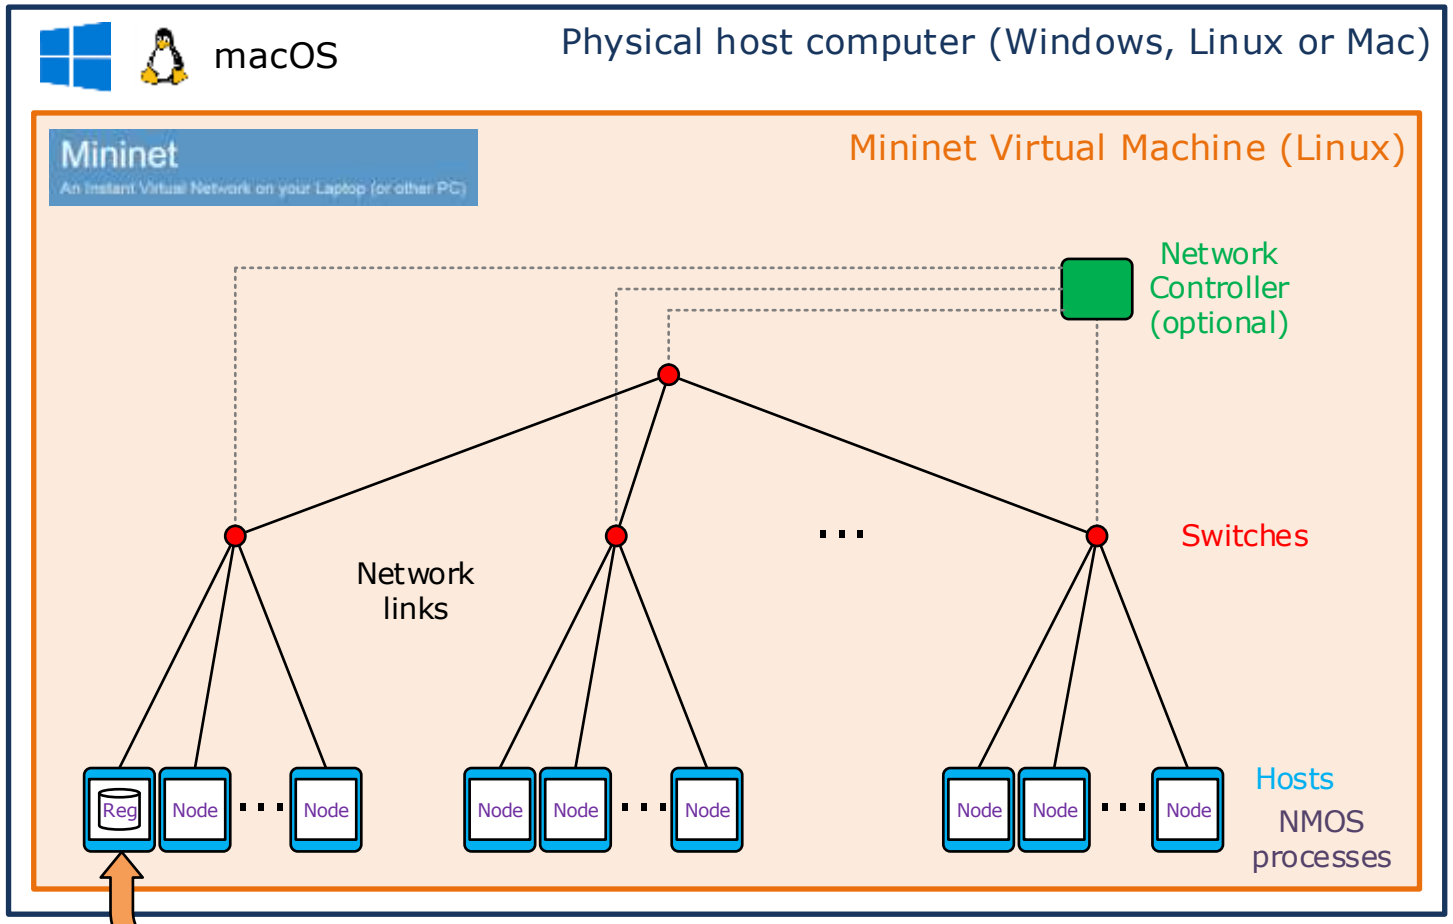

SONY

IP SHOWCASE™ THEATER

AMWA NMOS IS-04 and IS-05 APIs

The AMWA NMOS Scalability Study

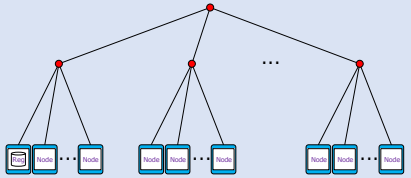

Scalability Study Results

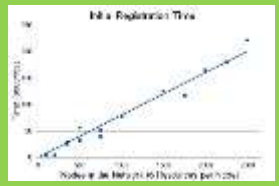

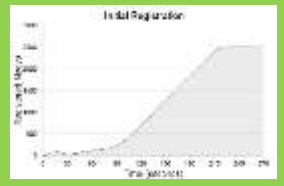

Best Practice Recommendations
and How JT-NM TR-1001-1 Helps

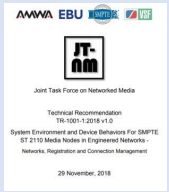

SONY

IP SHOWCASE™ THEATER

Initial registration

Test: Registration of 100 - 2,500 Nodes with 6 resources per Node - one of each resource type: Node, Device, Sender, Receiver, Source, Flow

Result: Total registration time scales linearly with number of resources

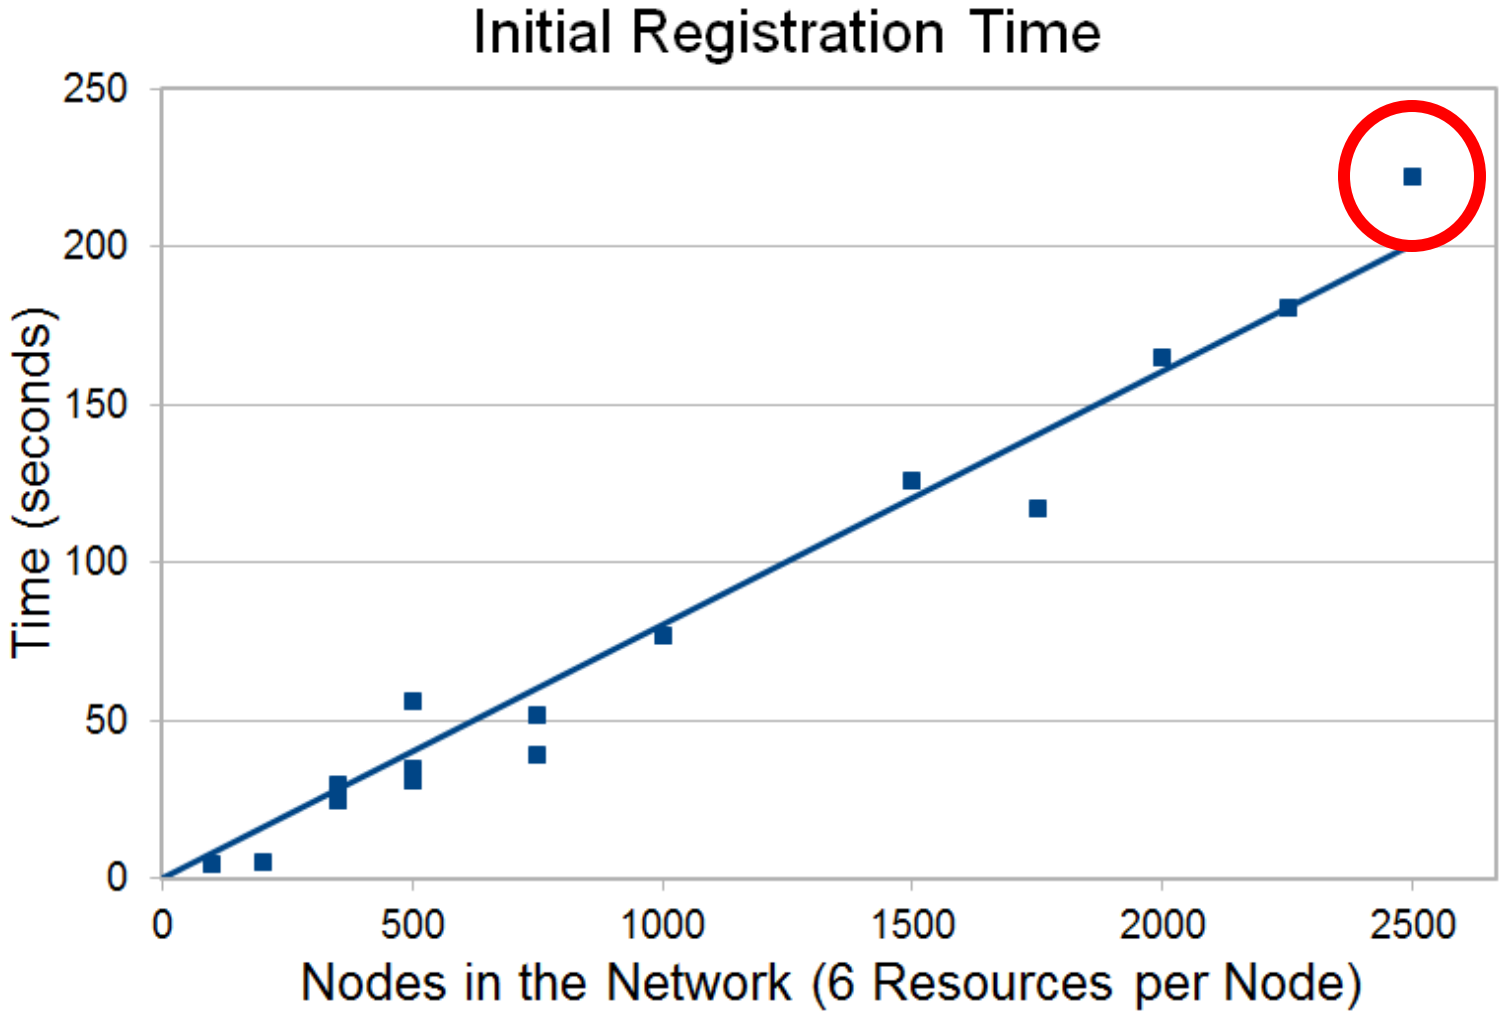

SONY

IP SHOWCASE™ THEATER

Initial registration

Test: Registration of 2,500 Nodes with 6 resources per Node

Result: Total registration time > 3.5 minutes

Rate of registration varies over time – slow start and long tail

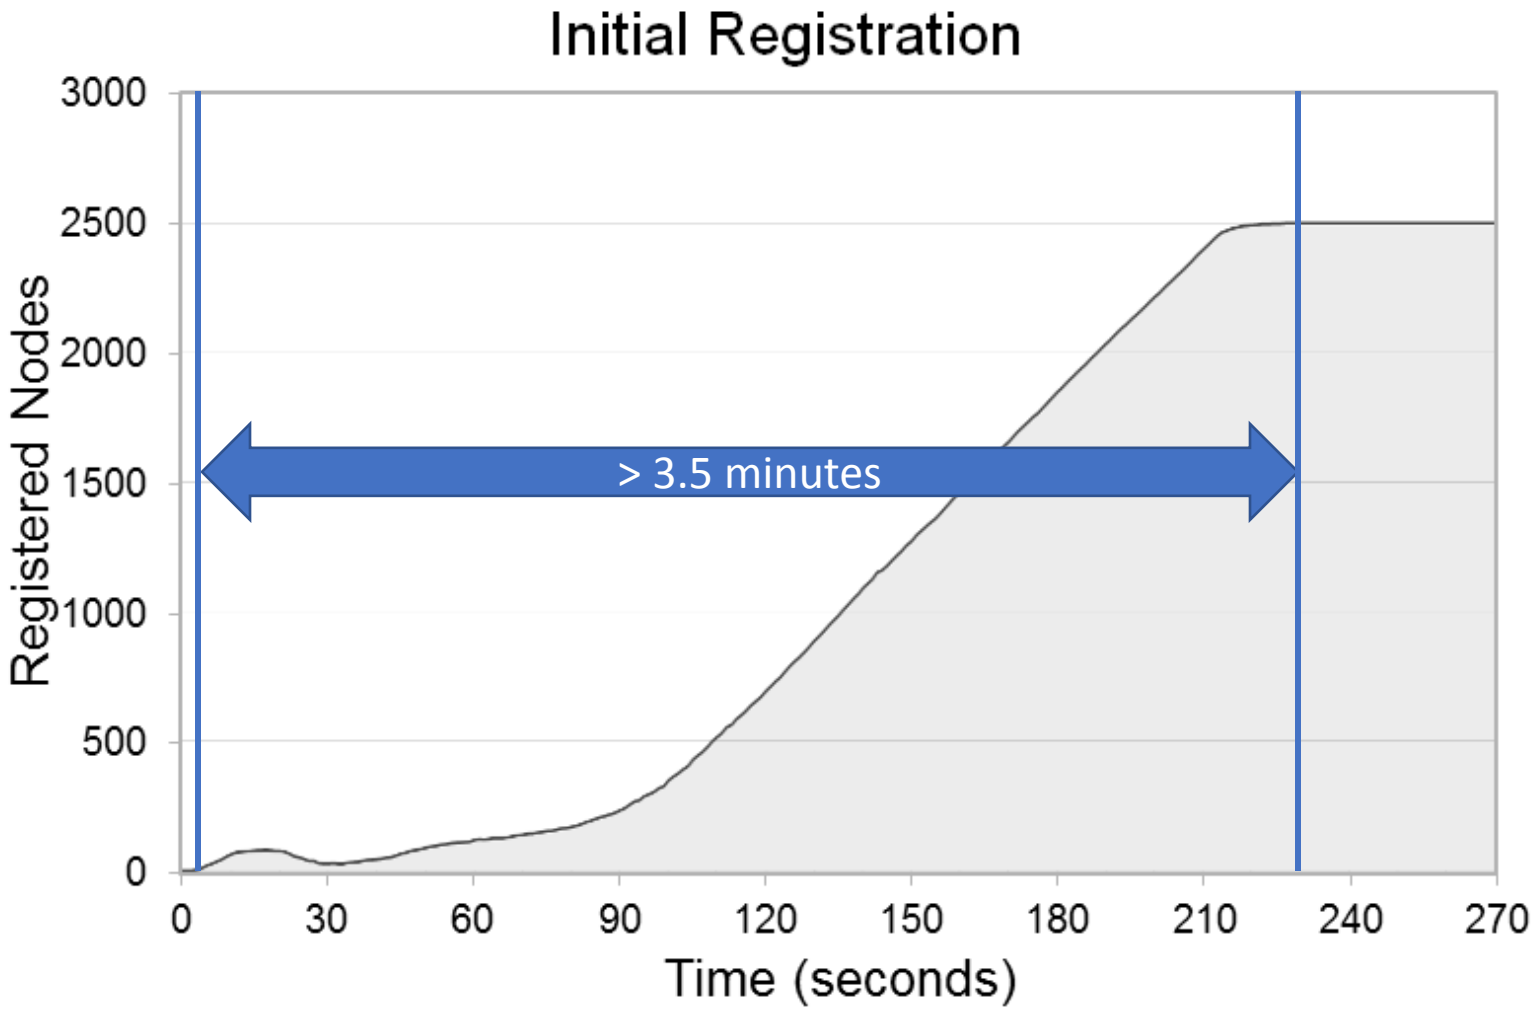

SONY

IP SHOWCASE™ THEATER

Initial registration

Test: Registration of 750 Nodes with 6 resources per Node

Result: Total registration time > 50 seconds

Slow start and long tail due to long DNS-SD and HTTP timeout and retry intervals

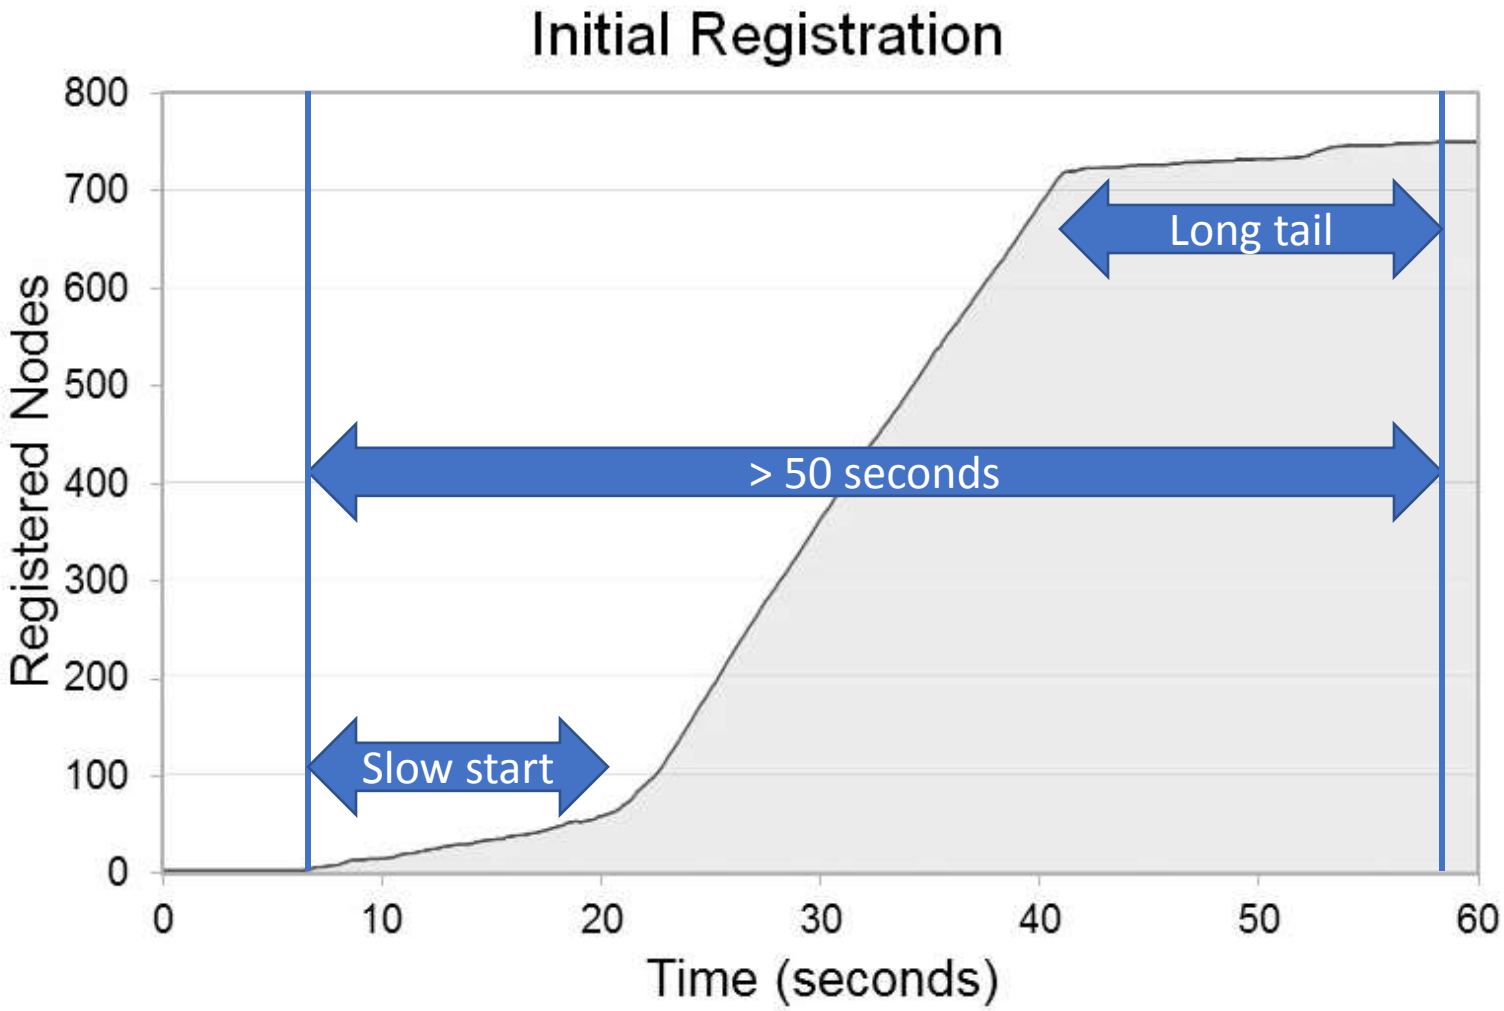

SONY

IP SHOWCASE™ THEATER

Initial registration

Test: Registration of 750 Nodes with 6 resources per Node.

With optimisations to:

- DNS-SD retry interval
- HTTP timeout

Result: Total registration time < 30 seconds

No slow start or long tail

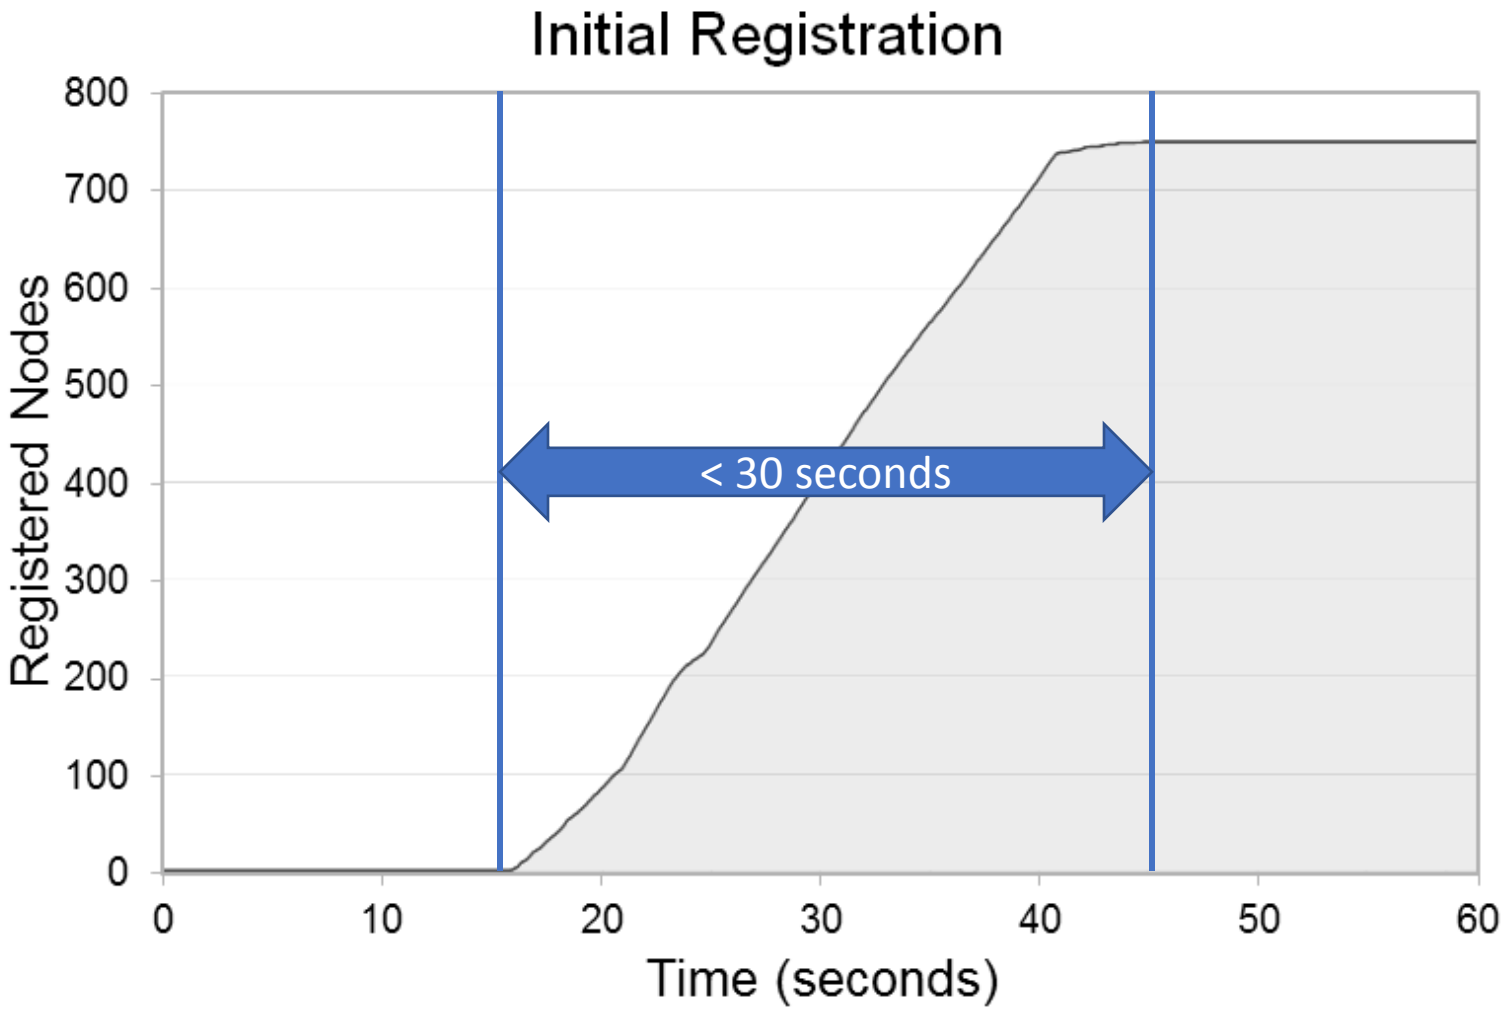

SONY

IP SHOWCASE™ THEATER

Recovery after a failure

Test: Register 750 Nodes with 6 resources per Node.

Break network link to registry for a short period

Result 1: Break for 5s -> Full recovery within 30s

Result 2: Break for 30s -> Full recovery within 90s

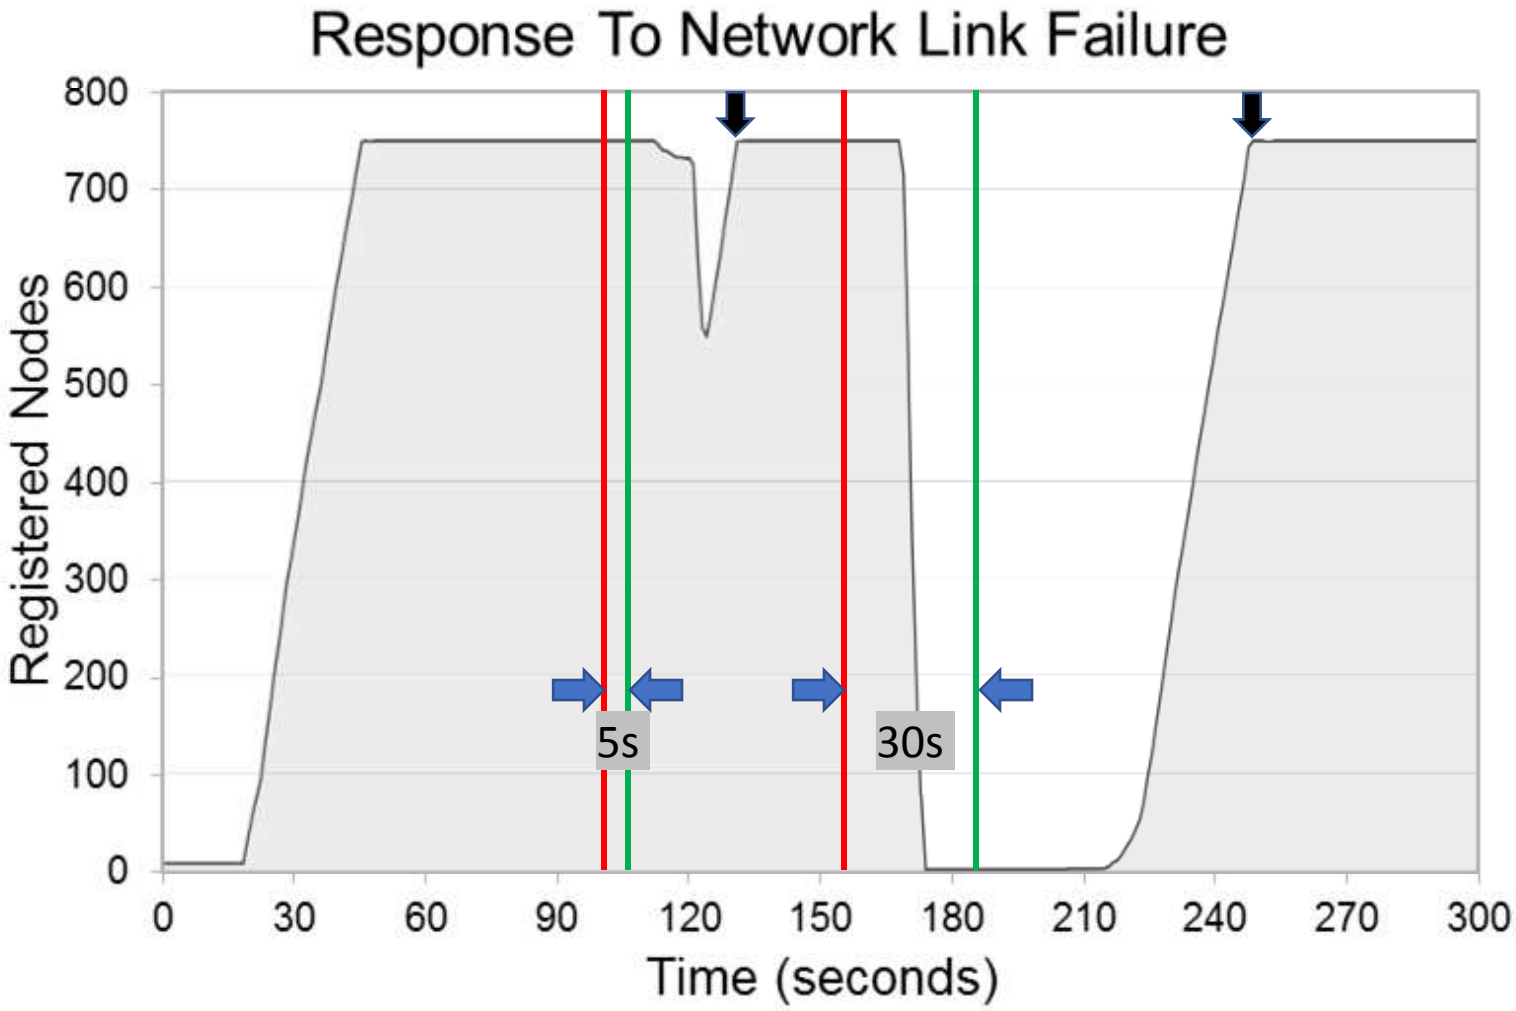

SONY

IP SHOWCASE™ THEATER

Multiple registries

Test: Registration of 750 Nodes with 6 resources per Node with two-way replication in a federated RDS

Result: Total registration time reduced from 30 seconds to <15 seconds

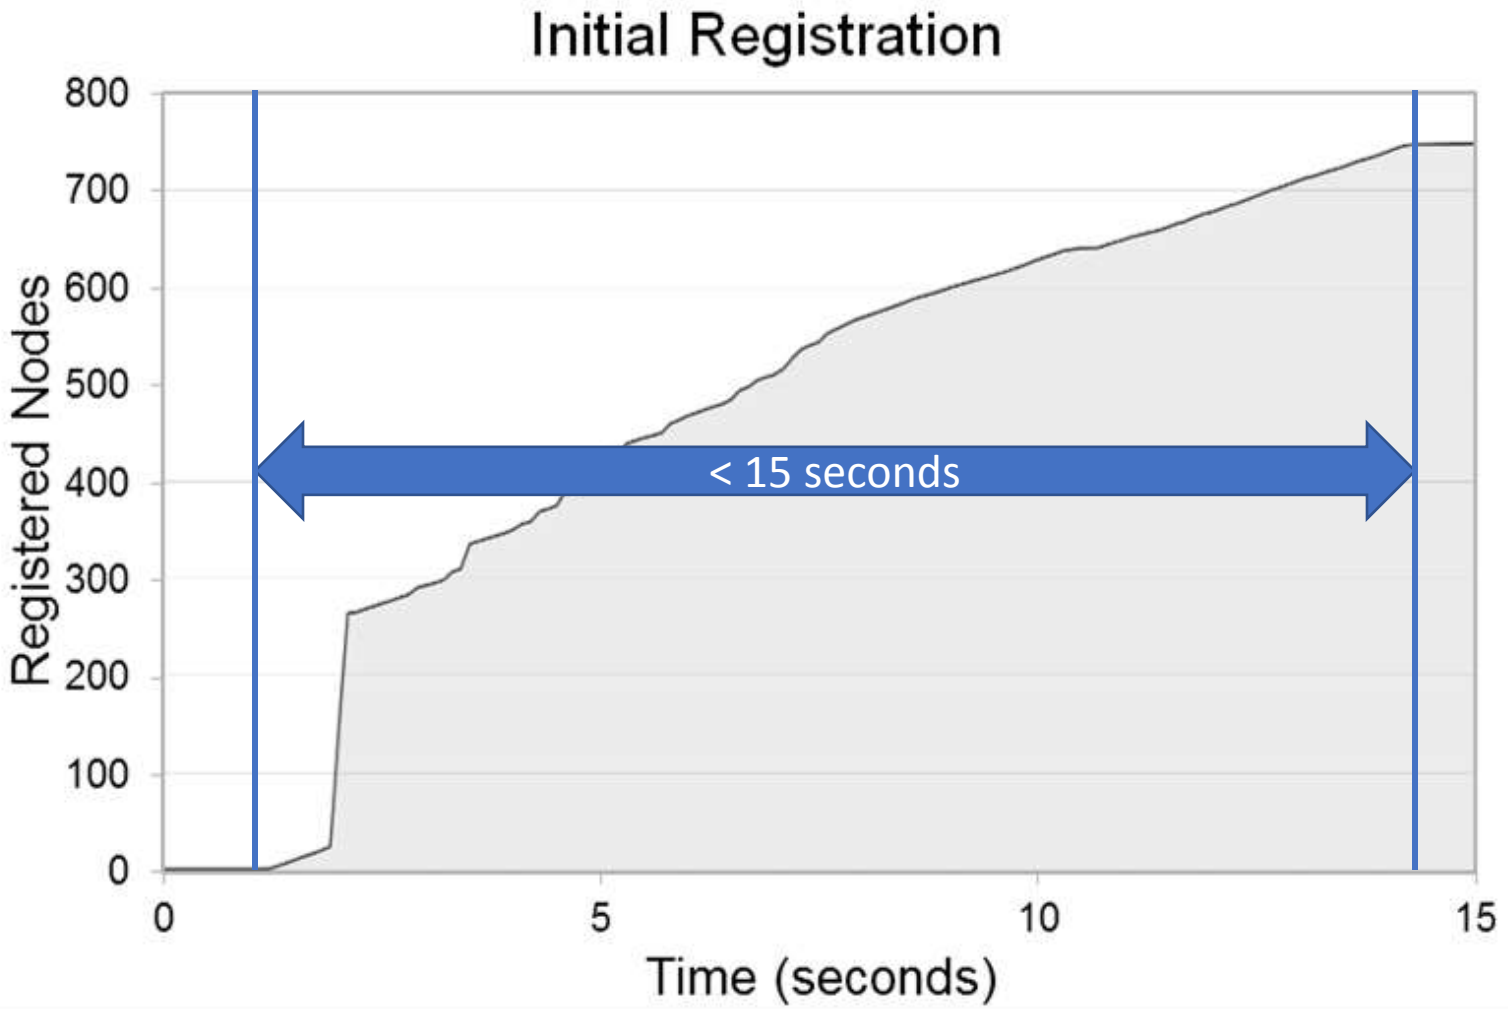

SONY

IP SHOWCASE™ THEATER

Connection management at scale

Test: Populate Client's crosspoint matrix (2,500 x 2,500) using IS-04 Query API

Result: Total time to populate matrix < 1.0s

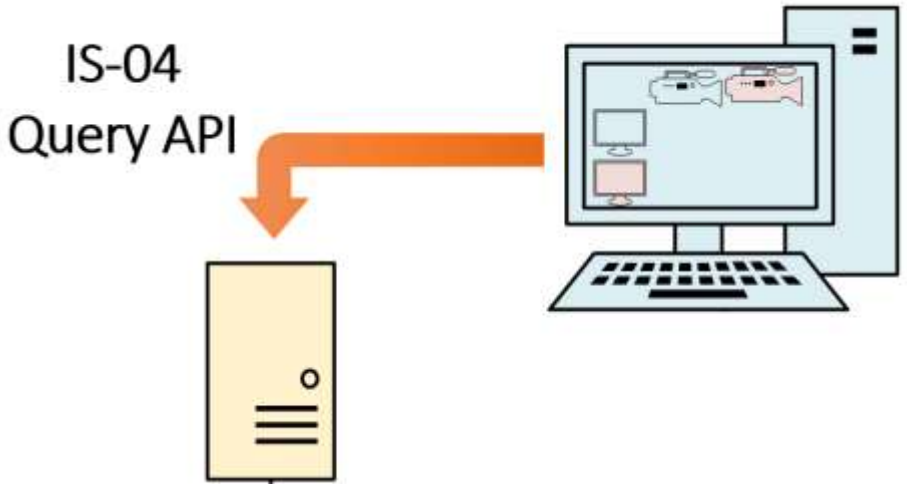

SONY

IP SHOWCASE™ THEATER

Connection management at scale

Test: Update 750 crosspoints using IS-05 Connection API

Result: Total time to update crosspoints < 3.0s

Optimised API usage is important - make good use of paging, push notifications and WebSocket API.

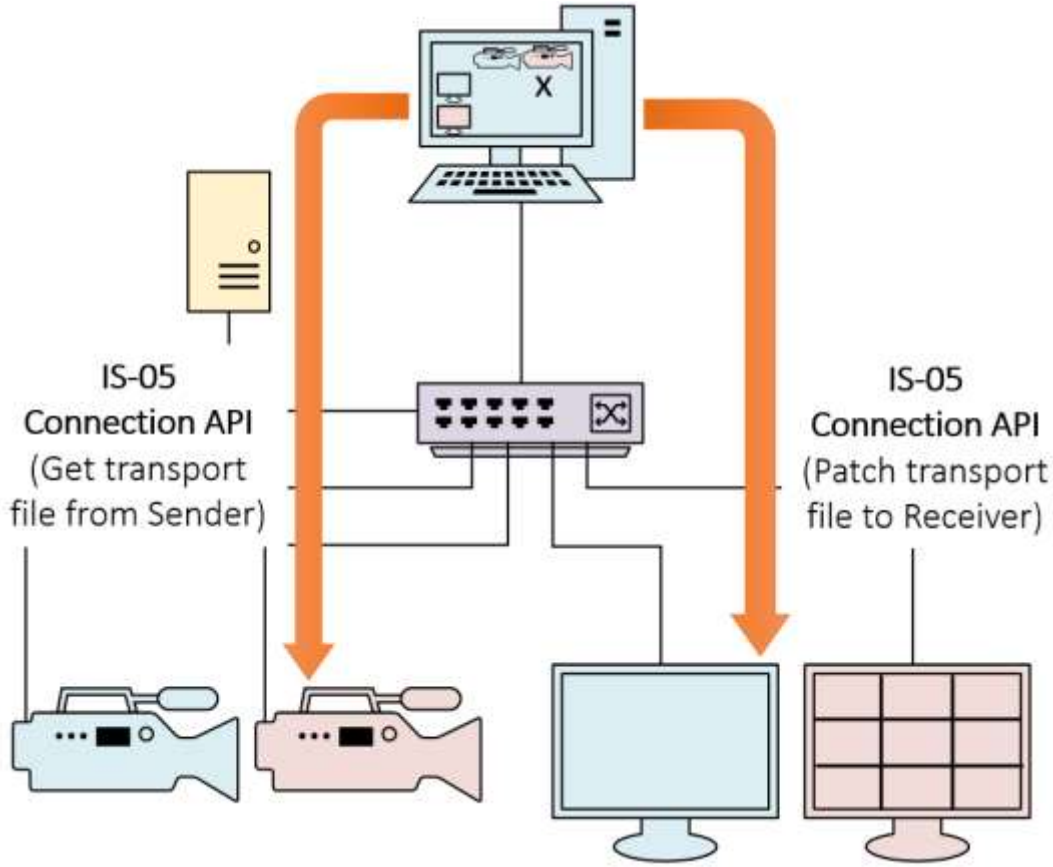

SONY

IP SHOWCASE™ THEATER

AMWA NMOS IS-04 and IS-05 APIs

The AMWA NMOS Scalability Study

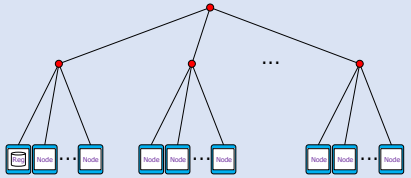

Scalability Study Results

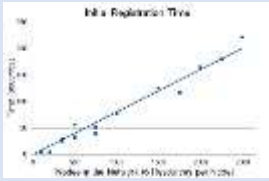

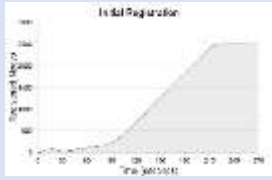

Best Practice Recommendations
and How JT-NM TR-1001-1 Helps

SONY

IP SHOWCASE™ THEATER

Best Practice Recommendation

- Use IS-04 and IS-05 across all Media Nodes
  - Ensures interoperability!

Media Nodes shall expose an AMWA NMOS IS-04 1.2 or higher node API, and shall register using the IS-04 registration API

Media Nodes shall expose an AMWA NMOS IS-05 1.0.2 STABLE or higher device connection management API

How JT-NM TR-1001-1 Helps

- Mandates use of IS-04 and IS-05

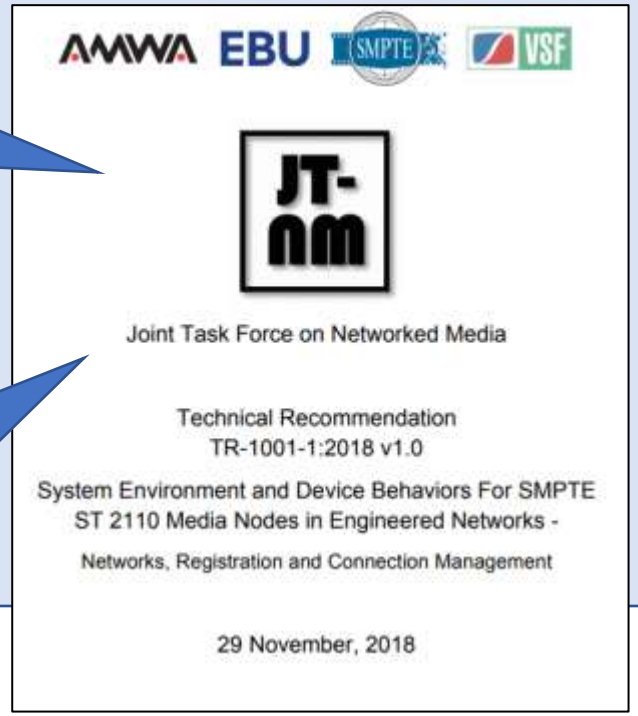
AMWA EBU SMPTE VSF

JT-nm

Joint Task Force on Networked Media

Technical Recommendation
TR-1001-1:2018 v1.0

System Environment and Device Behaviors For SMPTE ST 2110 Media Nodes in Engineered Networks - Networks, Registration and Connection Management

29 November, 2018

SONY

IP SHOWCASE™ THEATER

Best Practice Recommendation

- Use clustered / federated registries to improve performance
  - Faster registration
  - Better fault tolerance

The Network Environment shall contain **one or more registries** implementing AMWA NMOS IS-04 Registration and Query APIs

How JT-NM TR-1001-1 Helps

- Multiple registries supported

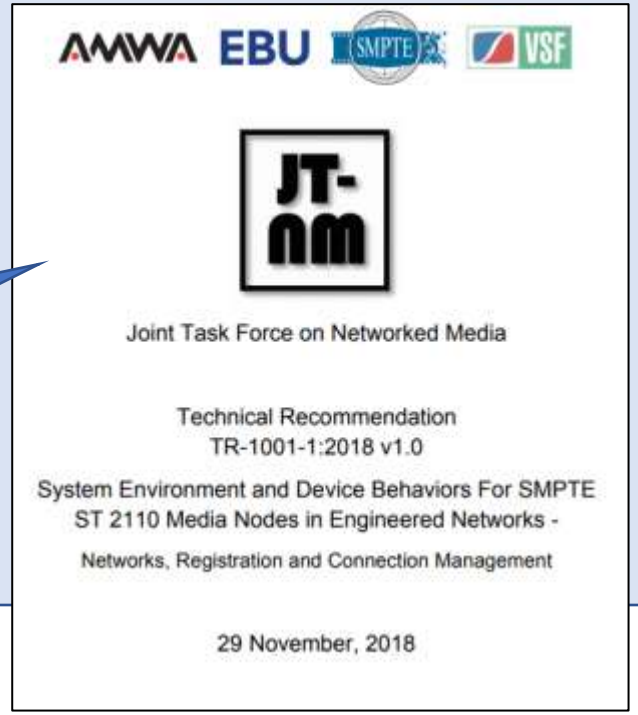

SONY

IP SHOWCASE™ THEATER

Best Practice Recommendation

- Choose heartbeat and registry expiry intervals carefully
  - Use common heartbeat intervals across all Media Nodes
  - Can be used to guide choice of DNS-SD and HTTP timeout and retry intervals
  - Recommended values from IS-04 spec worked well in our experiments

How JT-NM TR-1001-1 Helps

- Heartbeat interval is stored in System Resource and must be used by all Media Nodes

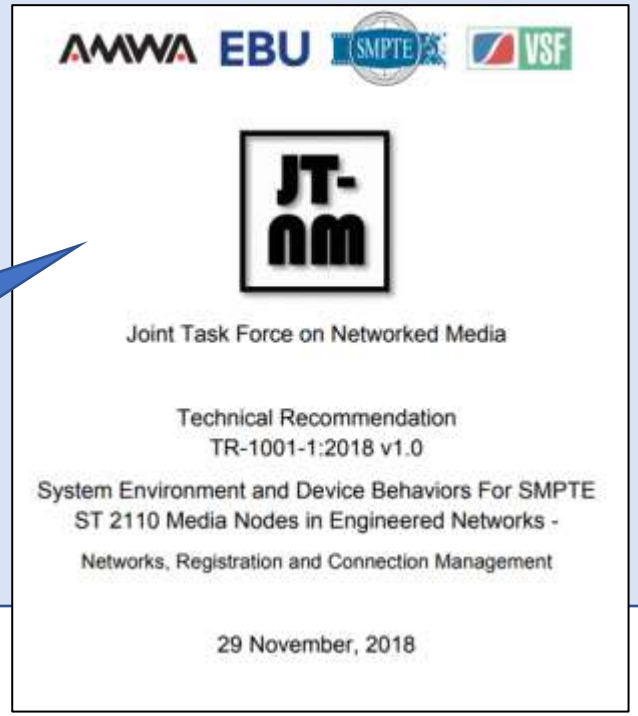
AMWA EBU SMPTE VSF

JT-nm

Joint Task Force on Networked Media

Technical Recommendation
TR-1001-1:2018 v1.0

System Environment and Device Behaviors For SMPTE ST 2110 Media Nodes in Engineered Networks - Networks, Registration and Connection Management

29 November, 2018

Media Nodes shall use the registry heartbeat_interval value specified in the System Resource … when maintaining their registration

SONY

IP SHOWCASE™ THEATER

Best Practice Recommendation

- Use unicast DNS-SD
  - Improved registration performance
  - Better for scalability
  - Essential for layer 3 networks
  - Multicast DNS is being deprecated from IS-04

Media Nodes shall use unicast DNS Service Discovery (DNS-SD) to locate the registration APIs as described in IS-04

How JT-NM TR-1001-1 Helps

- Unicast DNS-SD is mandated

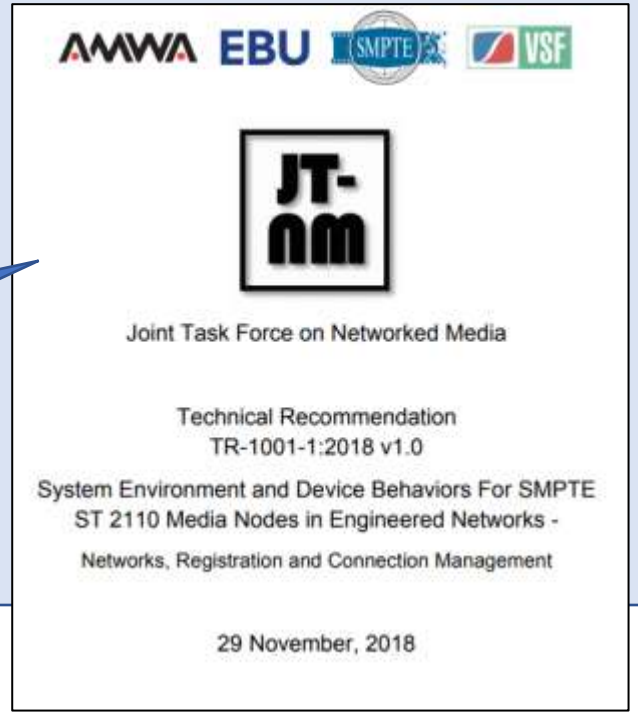
AMWA EBU SMPTE VSF

JT-nm

Joint Task Force on Networked Media

Technical Recommendation
TR-1001-1:2018 v1.0

System Environment and Device Behaviors For SMPTE ST 2110 Media Nodes in Engineered Networks - Networks, Registration and Connection Management

29 November, 2018

SONY

IP SHOWCASE™ THEATER

Resources

- AMWA NMOS Scalability Study test environment
  - https://github.com/AMWA-TV/nmos-scalability
- Sony nmos-cpp open source software for IS-04 Registry and IS-04/-05 Node
  - https://github.com/sony/nmos-cpp
- SMPTE Annual Technical Conference 2018 paper
  - “Scalability and Performance of the AMWA IS-04 and IS-05 NMOS Specifications for Networked Media” – Robert Porter and Gareth Sylvester-Bradley
  - https://ieeexplore.ieee.org/document/8610041

IP SHOWCASE™ THEATER

CURATED BY VSF VIDEO SERVICES FORUM

Thank You

Rob Porter, Sony Europe B.V.

Rob.Porter@sony.com