Performance Verification Architecture Task Force

Update

Jerry Sobieski (NORDUnet)

Oct 4, 2013

Singapore

NORDUnet

Nordic infrastructure for Research & Education

Performance Verification Architecture

- Objective (from the charter):
  - The GLIF End-to-End Performance Verification Architectures task force is chartered to develop recommendations for a deterministic, scalable, and secure architecture for determining the delivered end to end performance characteristics of emerging light path (connection oriented) network services.
- Co-Chairs
  - Steve Wolff (Internet2)
  - Jerry Sobieski (NORDUnet)

Progress since Honolulu

- Discussion began in early 2012…
- Limited participation… (still not sure why…)
- Basic conference call: "Why can't we just use perfSONAR?..."
- …need a paper describing future services, why verification is important in this future environment, what needs to be verified, when, where, how by whom,… What capabilities will a comprehensive and deterministic PV process require?
- Paper is [almost] available:
- Purpose: Explain why verification of performance *guarantees* for Connection Services is different from characterizing conventional best effort performance – and why deterministic methods are necessary not just to insure delivery of predictable performance, but to understand how such services behave – or fail.

  → Get GLIF community to recognize why PV is critical for the success of guaranteed services.
- **"Performance Verification Architecture and Emerging Performance Guaranteed Network Services"**
  - …some editing is still in progress (18 pages + illustrations).
  - Will be circulated for review before Jan 2013.
- Summary follows…

Progress since Honolulu

(This space unintentionally left blank.)

Where to from here?

- PV is important
  - If we do not or cannot verify – somehow - our service as a user or as a provider, then we cannot rely upon them and guarantees become pointless.
- We *need* predictable services
    - Capacity/performance, path routing, reliability, security, scheduling, ...
  - Thus we need service guarantees
  - Thus we need a PVA that can determine is a service instance is conformant
- GLIF should not write off PVA simply due to unavailable manpower or prioritization that prevents progress
  - But perhaps we can approach it differently...

Forward...

- GLIF can de-commission the PVA Task Force ...for now.
  - We can always re-activate later if/when we have manpower and tangible interest
  - Or we can give it another 6 months and try [harder] to make progress... (idle doesn't cost anything...)
- GLIF should solicit talks and research for topics around these issues:
    - Performance Verification
    - PV at 1Tbps
    - Performance tuning and measurement for Isochronous configuration (packet buffer configuration for performance guarantees)
    - Automated Fault Localization
    - Inter-domain PV – both for circuits and best effort packet performance

Refresher

- The following slides pose one scenario for Performance Verification
  - Food for thought, experiments, research proposals, ...

Deterministic End to End Performance Verification Architecture

Offered for discussion to the GLIF Performance Verification Architectures Task Force

Jerry Sobieski
NORDUnet
Oct 4, 2013
Singapore

NORDUnet
Nordic infrastructure for Research & Education

Performance

"Measure what can be measured, and make measurable what cannot be measured."

- Galileo Galilei

- Definition: "Performance" is a possibly multivariate quantity that can be measured to an agreed-upon precision by an agreed-upon measurement protocol.

The Problem

- Emerging Connection Services offer “guaranteed” performance ……or so they say.
- How do we verify this performance?
  1. Determining when a Connection is performing as requested/required…or not.
  2. Determining which aspects of the performance guarantees are not functioning to spec
  3. Determine (to some resolution) “where” a Connection is failing
- Guaranteed services are intended to provide deterministic service – predictable, reliable, repeatable… And so require substantially tighter engineering constraints than best effort
  - Deterministic PV processes are critical
- If performance is flawed, it needs to be fixed, ASAP.
  - Identifying the under-performing segment and notifying the agent in charge via automated means is End Game (Automated Fault Localization)

What are we “verifying”?

- What networks do: transport user data from one point to another.
- Performance guarantees are network services that make that transport behavior *predictable* and *repeatable*.
- Connections are “logical conduits that carry user data transparently and unmodified from an ingress location to an egress location” (ref: OGF NSI Framework 2011)

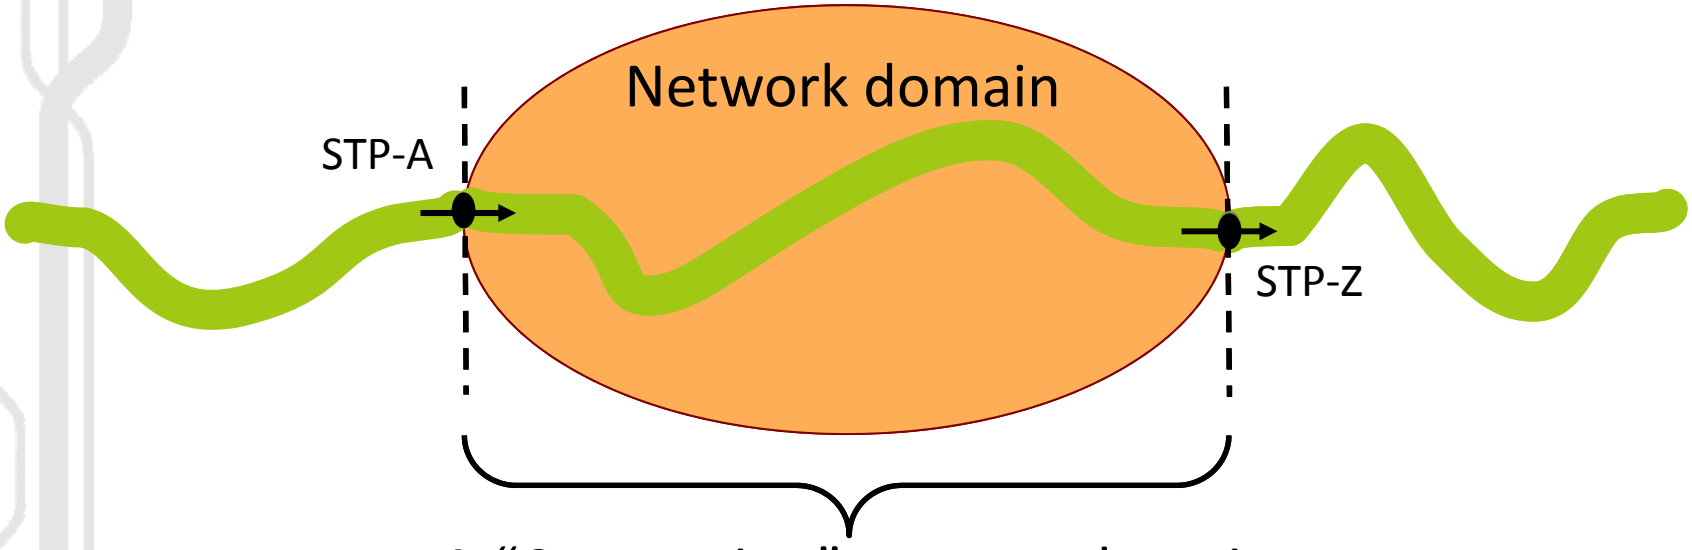

NORDUnet
Nordic infrastructure for Research & Education

Global *inter-Domain* Services

- In the real world…End to End connection services are predicated across multiple network service domains…
- "Inter-operating common services domains" are not equivalent to "a single administrative domain"…
  - These are separate and independent organizations and services
  - With local realities: levels of expertise, performance capabilities, authorization constraints, security/privacy concerns, accounting, legal environments, peering arrangements,…

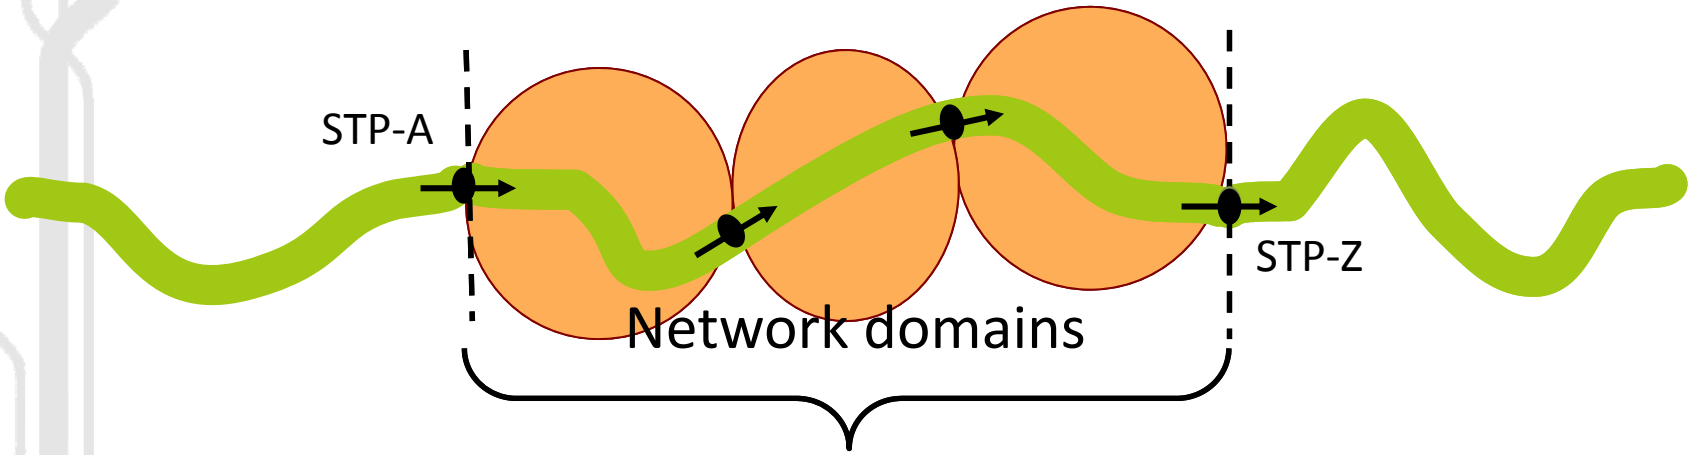

How do we verify the performance end to end?

Why verify?

- As Users, we do not trust the provider.
  - The provider may be incompetent
  - The provider is a shyster – the PA knowingly does not enforce the performance guarantees ("we have 100 Gbps core – we don't need to enforce bandwidth constraints in the core...")
  - The service itself is ill-defined wrt certain performance aspects
  - s#!t happens – something outside of the control of the provider broke.
- The user needs an means of determining whether the service they received meets the specifications they requested and were promised
  - That mechanism must be independent of the [untrusted] provider.

Why verify?

- As Providers, we do not trust the user:
  - Users are incompetent – they do not configure their end system(s) or access networks to properly
  - The user is not conforming to the service guarantees (or the service is undefined with respect to some aspect.)
  - s%#$@t happens – something outside of our control broke – we can not rely on the user to notify us prior to the law suit.
- The Provider needs an means of determining whether the service they delivered meets the specifications they guaranteed
  - Anf that mechanism must be independent of the [untrusted] user

Independent Verification

- "Verification" means corroborating what the other party claims to be true.
- "Independent verification"
  - The user's verification testing does not rely on their provider's assets/agents to provide corroboration
  - And the provider's tests likewise do not rely on the user's assets or agents for its results.
- The PV architecture must be able to provide an independent testing/measurement model.

Deterministic Performance Verification

- Can we *deterministically* measure the performance of a Connection?
- Can we do so without perturbing the flow?
- Can we do so in such a fashion that we can determine where along the path performance problems are occurring?

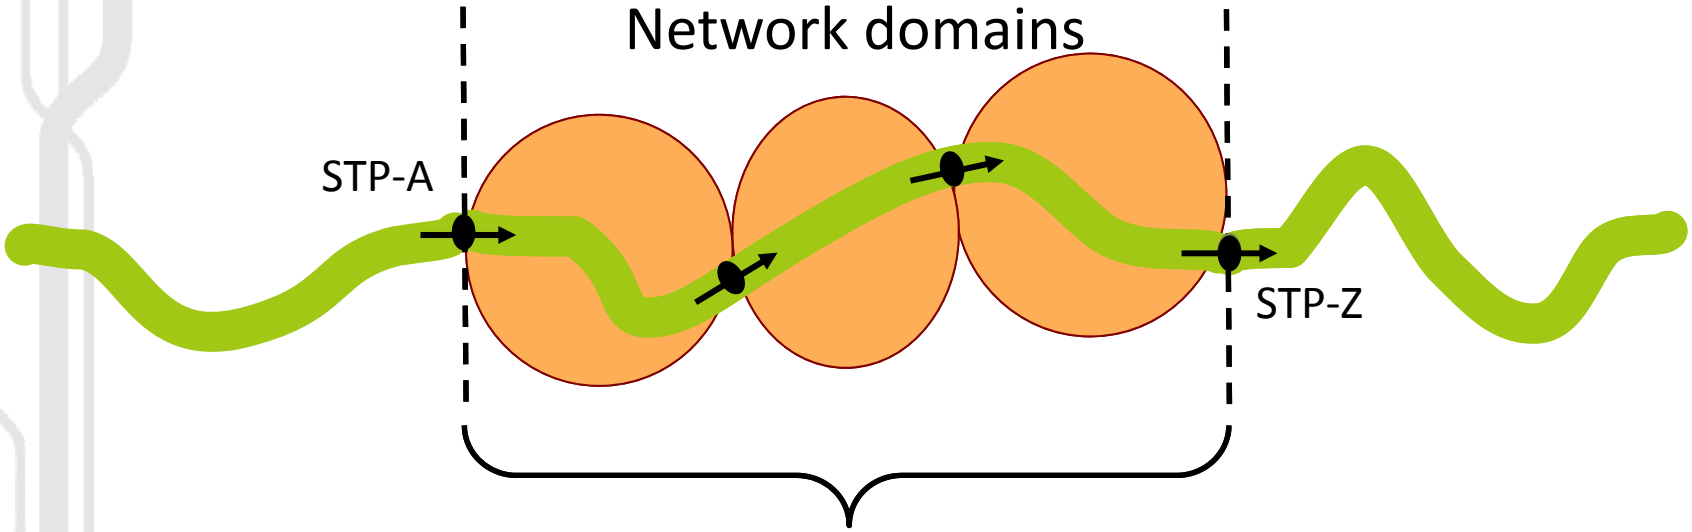

How do we verify the throughput from STP-A to STP-Z?

The Service Provider Demarc

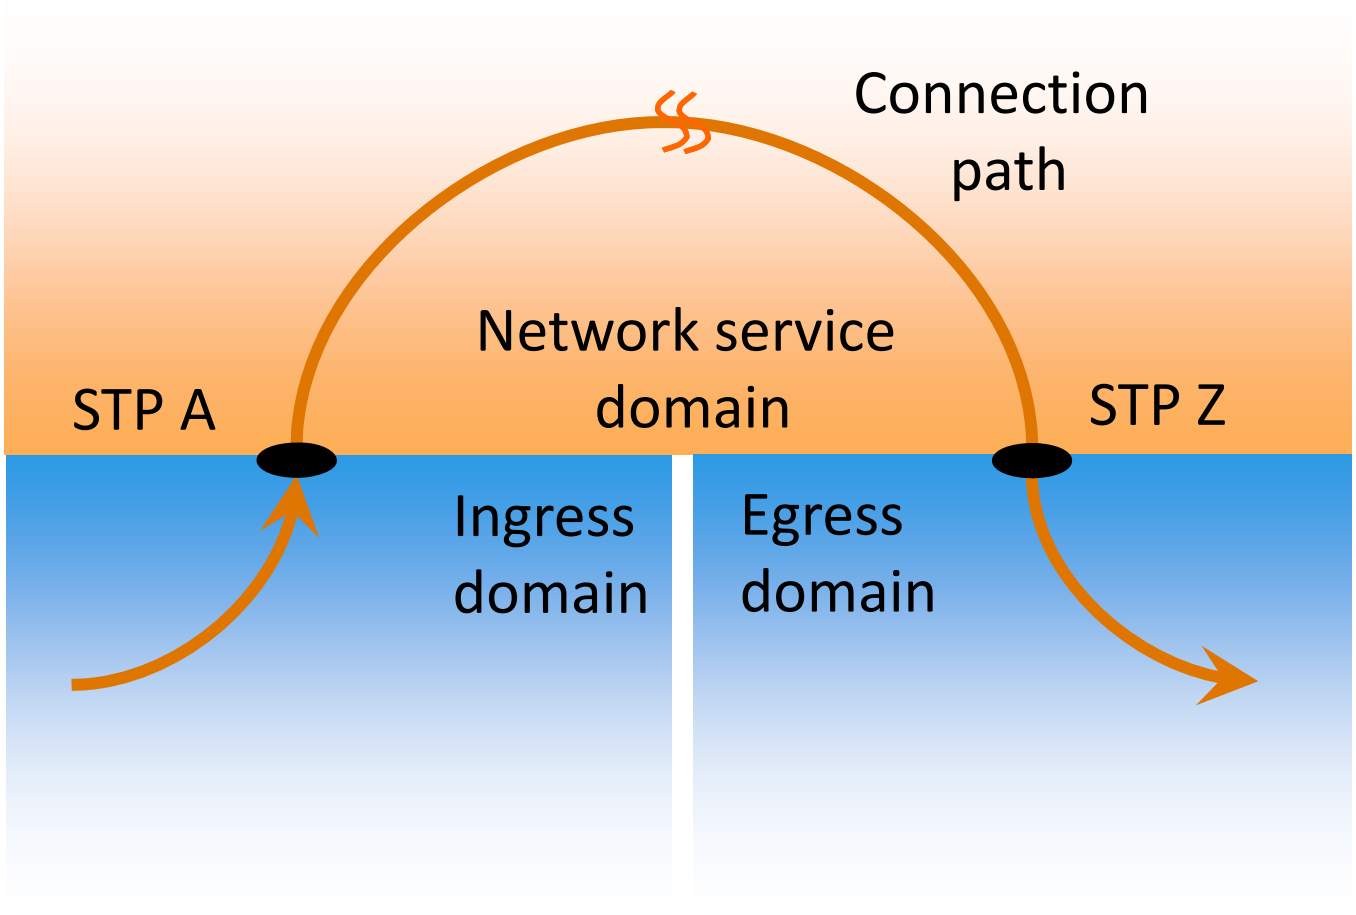

NORDUnet
Nordic infrastructure for Research & Education

PV Test Point Engineering

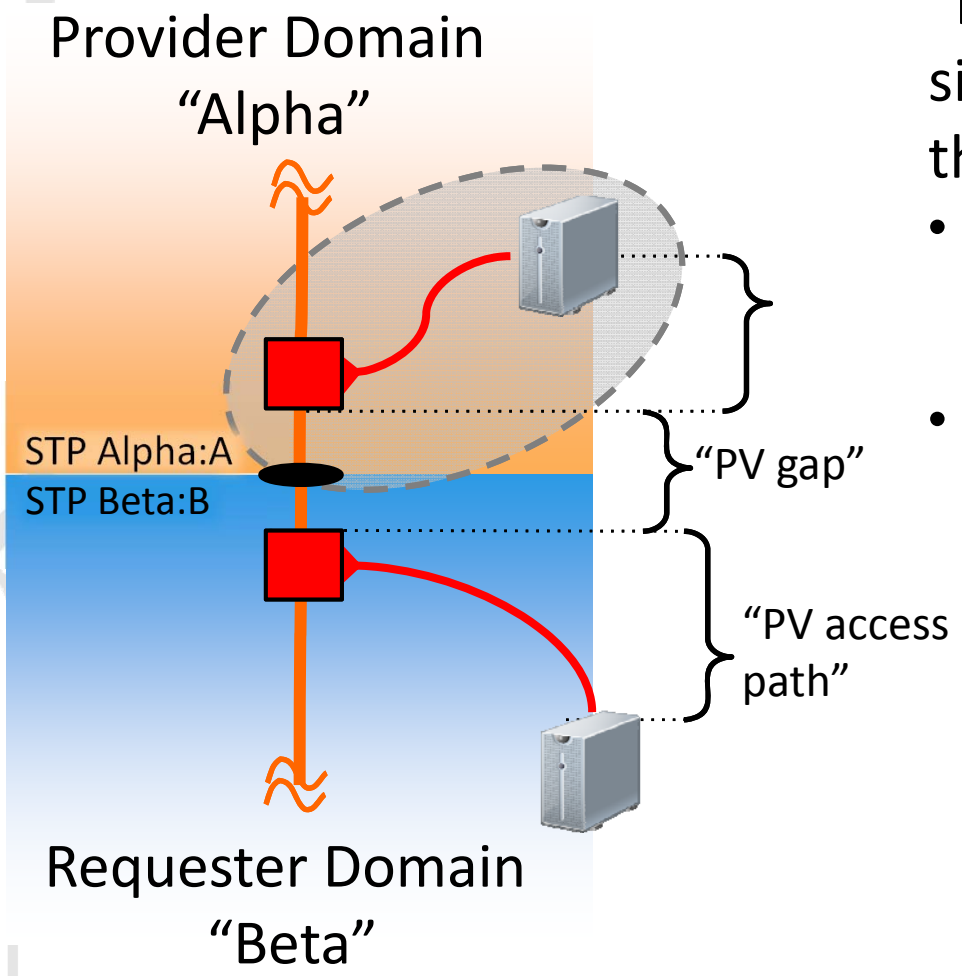

"PV Test Points" must be: situated as physically close to the demarc STP as possible

- Minimize effects of difference ("PV gap") between PV test points in adjacent domains
- Compact PV access in order to minimize path effects introduced by PV capture path itself

NORDUnet
Nordic infrastructure for Research & Education

PV Test Point Infrastructure

- "Tap" – A device in the data plane path that allows traffic to be inserted or observed.
- "PV Server" – an intelligent device that can source or sink data flows via the PV tap.
- PV Interface – the data channel between the PV Tap and the PV Server. (Note this may not be a simple patch cable)
- This Tap is a logical architectural feature since some transport infrastructure may be virtual

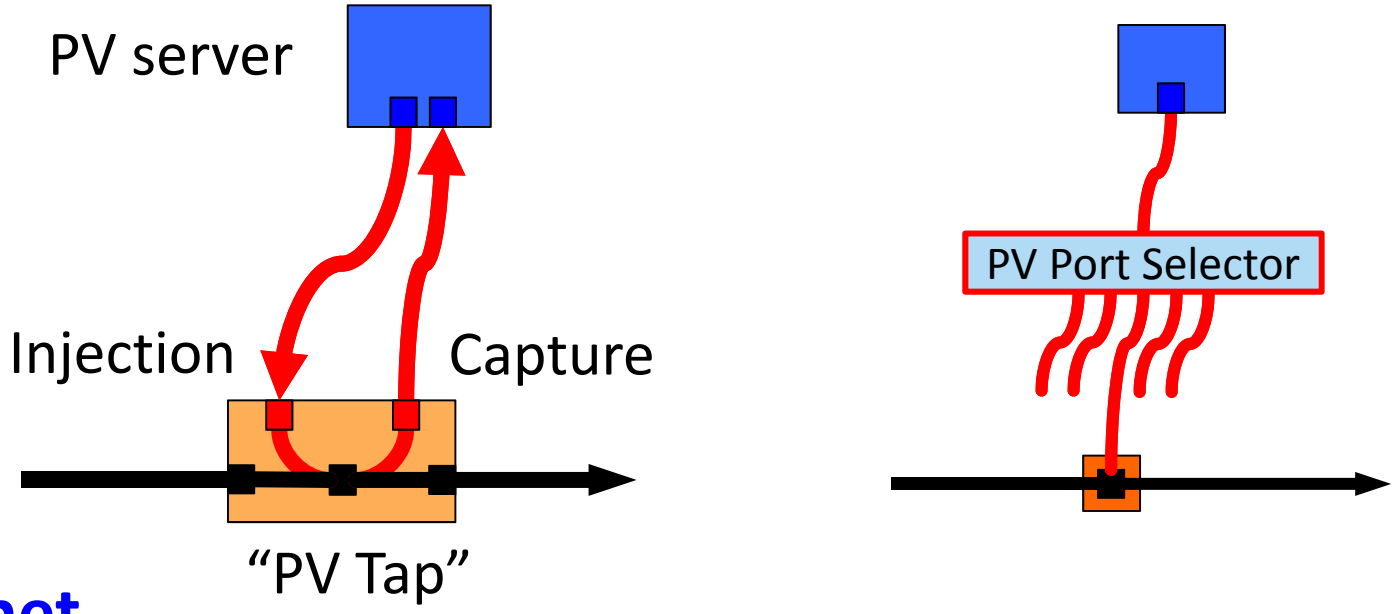

PV “Test Scope”

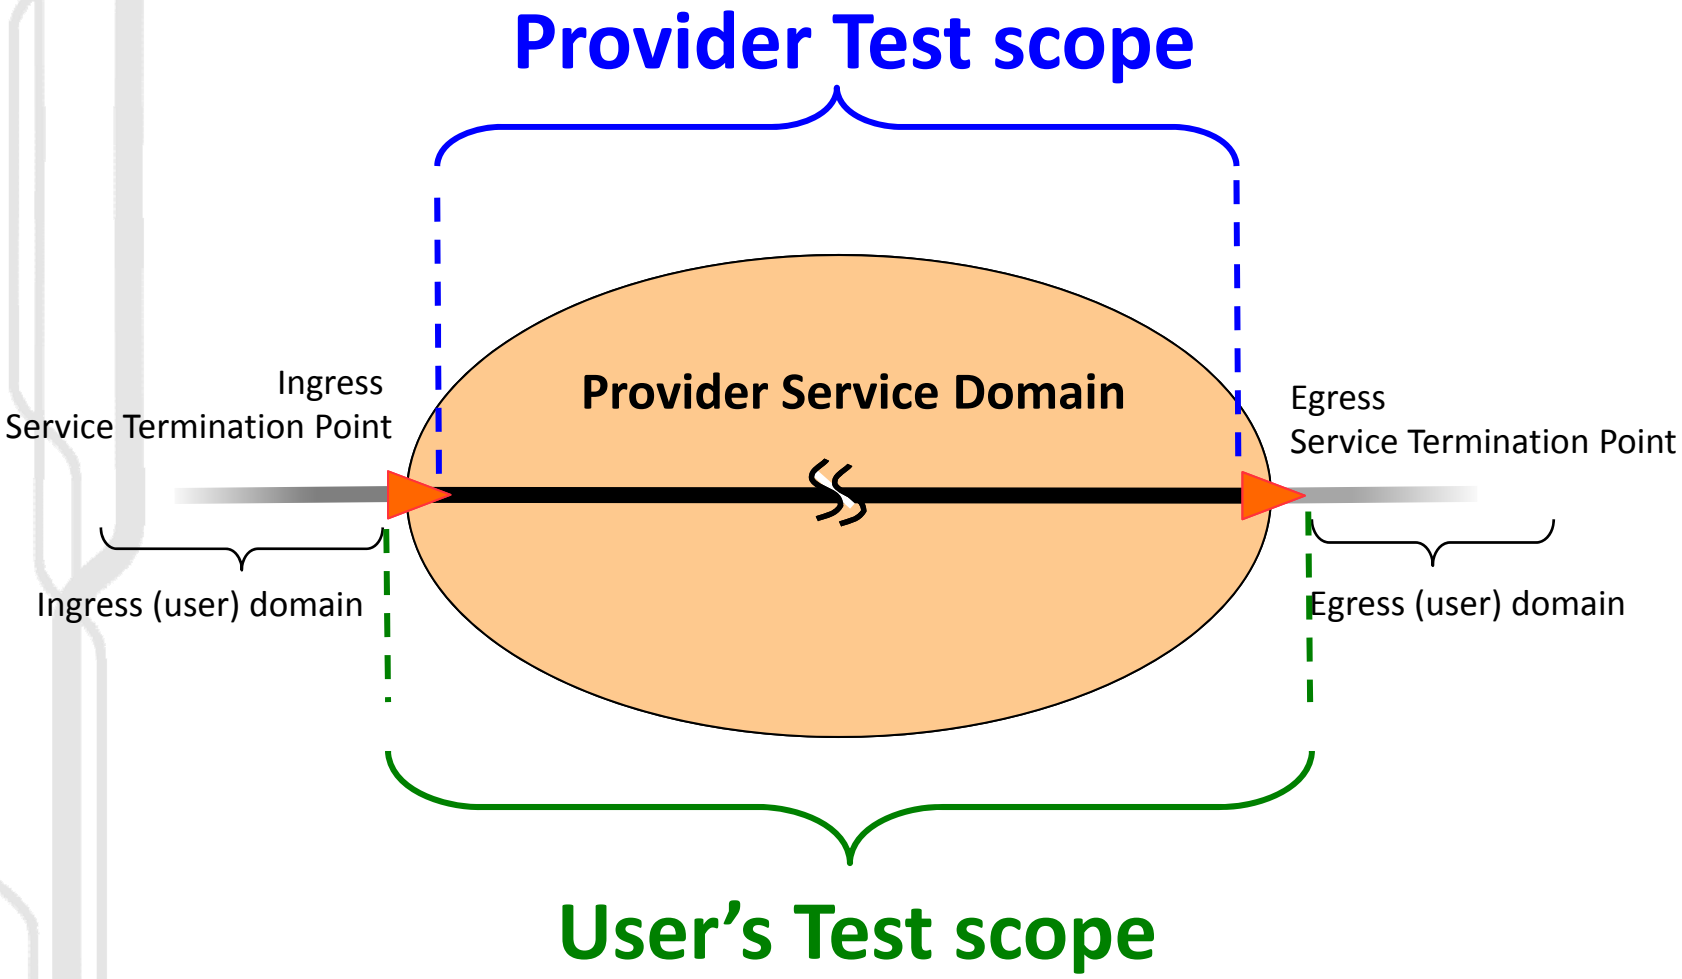

NORDUnet

Nordic infrastructure for Research & Education

Performance Flow Correlation

- Understanding the behaviour of jitter, latency, and good put of a guaranteed service ("connection") requires comparing the observed egress flow to the known ingress flow characteristics – "*flow correlation*"
- Deterministic PV captures and characterizes the source flow as tightly as the destination flow, and then the two flows are compared.
  - The comparison requires one or both captured flows be transferred to another host for correlation.
- Near real-time results can be had for most tests (delays of only seconds or minutes for large flow analysis.)
  - Flow correlation processing (transferring ingress and egress flows) must not interfering with test flows themselves.

Passive Performance Verification

- Passive testing is preferred
- Passive PV is determined based upon observed traffic…
  - An independent source could send artificial (and potentially non-representative) traffic profiles
  - Or the test points can observe actual user traffic
- Does not require taking circuit out of service
  - The ingress test point and egress test points just listen (and capture)
- Can be run at any time(!) without affecting the user's actual flow(!!)
  - Provider can periodically perform a passive flow correlation to verify performance while in service
  - Provider can continuously monitor/capture flows in order to observe or replay specific failure events

    (Probably unrealistic to record large flows over long periods – but circular buffers can provide "cockpit recorder" forensics.)

Active Performance Verification

- The active PV Server must be able to generate appropriate test flows:
  - Capacity – 100 Gbps
  - Accurate shaping capabilities
- Some aspects of performance guarantees cannot be verified except in the presence of resource contention (competing congestive flows)
  - PV server(s) must be able to source multiple flows (at high capacity and with accurate shaping)
  - Servers must be able to coordinate/synchronize flows to create exterior conditions that meet particular verification requirements

A “Performance Flow Correlator”

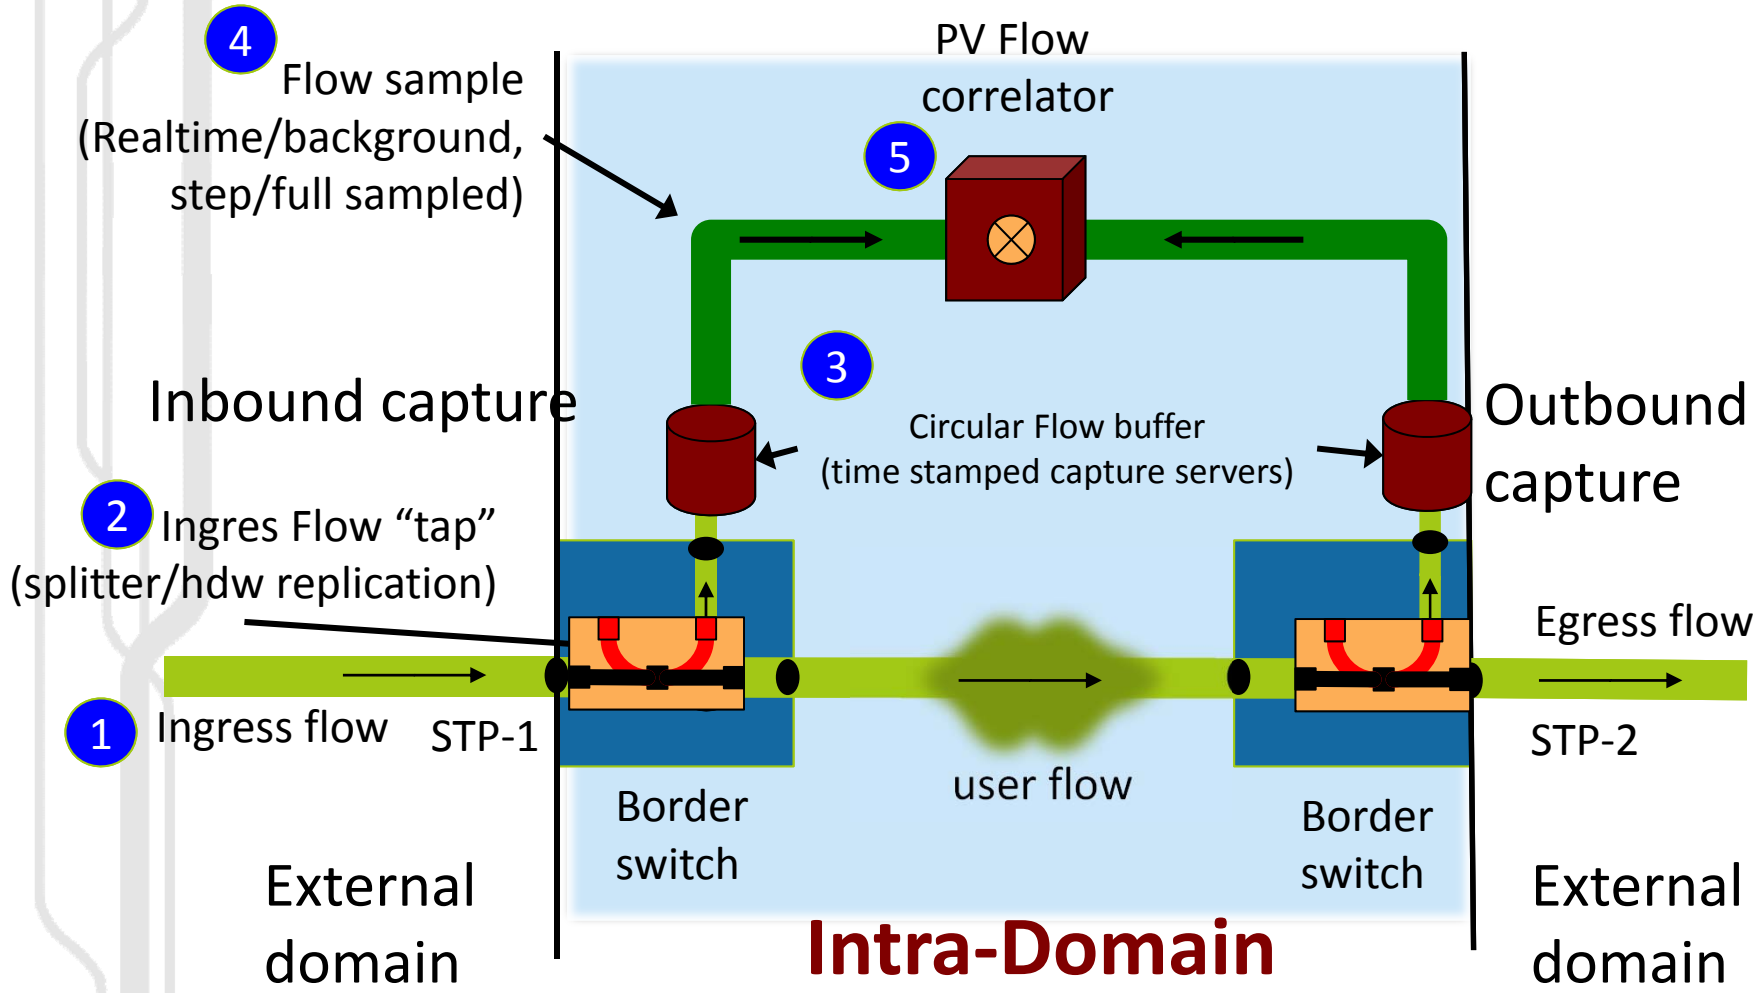

Intra-Domain Passive Provider PV Flow Correlation

NORDUnet
Nordic infrastructure for Research & Education

How the Flow Correlator Works

- The “PV tap” is implemented at every domain boundary interface
- The tap, when enabled, leads directly to a local capture device.
- The capture device must be able to:
  - Timestamp each datagram (depending on protocol).
  - Spool the captured stream to long term storage at line rate.
- The PV Capture Server can be sized and configured to capture an entire flow, or it can be sized to sample flows according to some rule or policy.
  - The correlation can be done in real time if engineered to do so…
  - or the flow can be captured and stored for later background analysis. … a few seconds later, or a few days later.
  - Correlation can be performed periodically, using short samples or real flows

An inter-Domain "Flow Correlator"

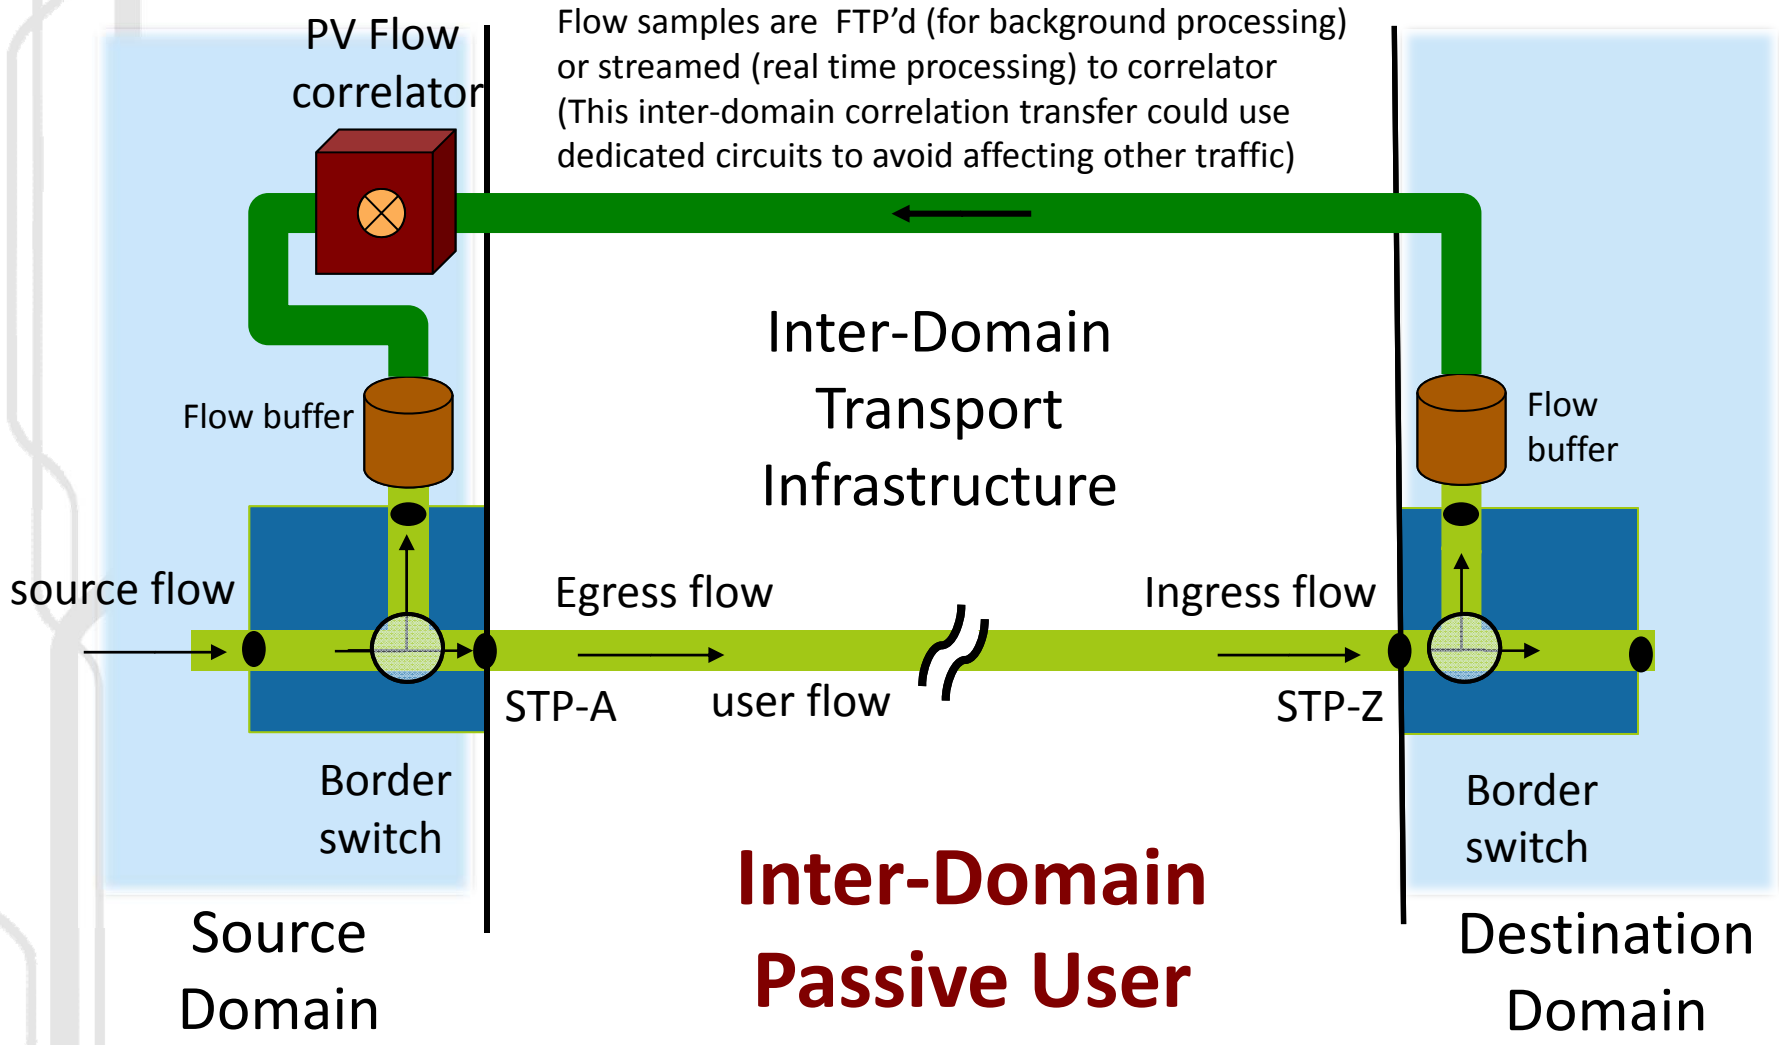

Inter-Domain Passive User Flow Correlation

NORDUnet
Nordic infrastructure for Research & Education

End to End PV as a Service

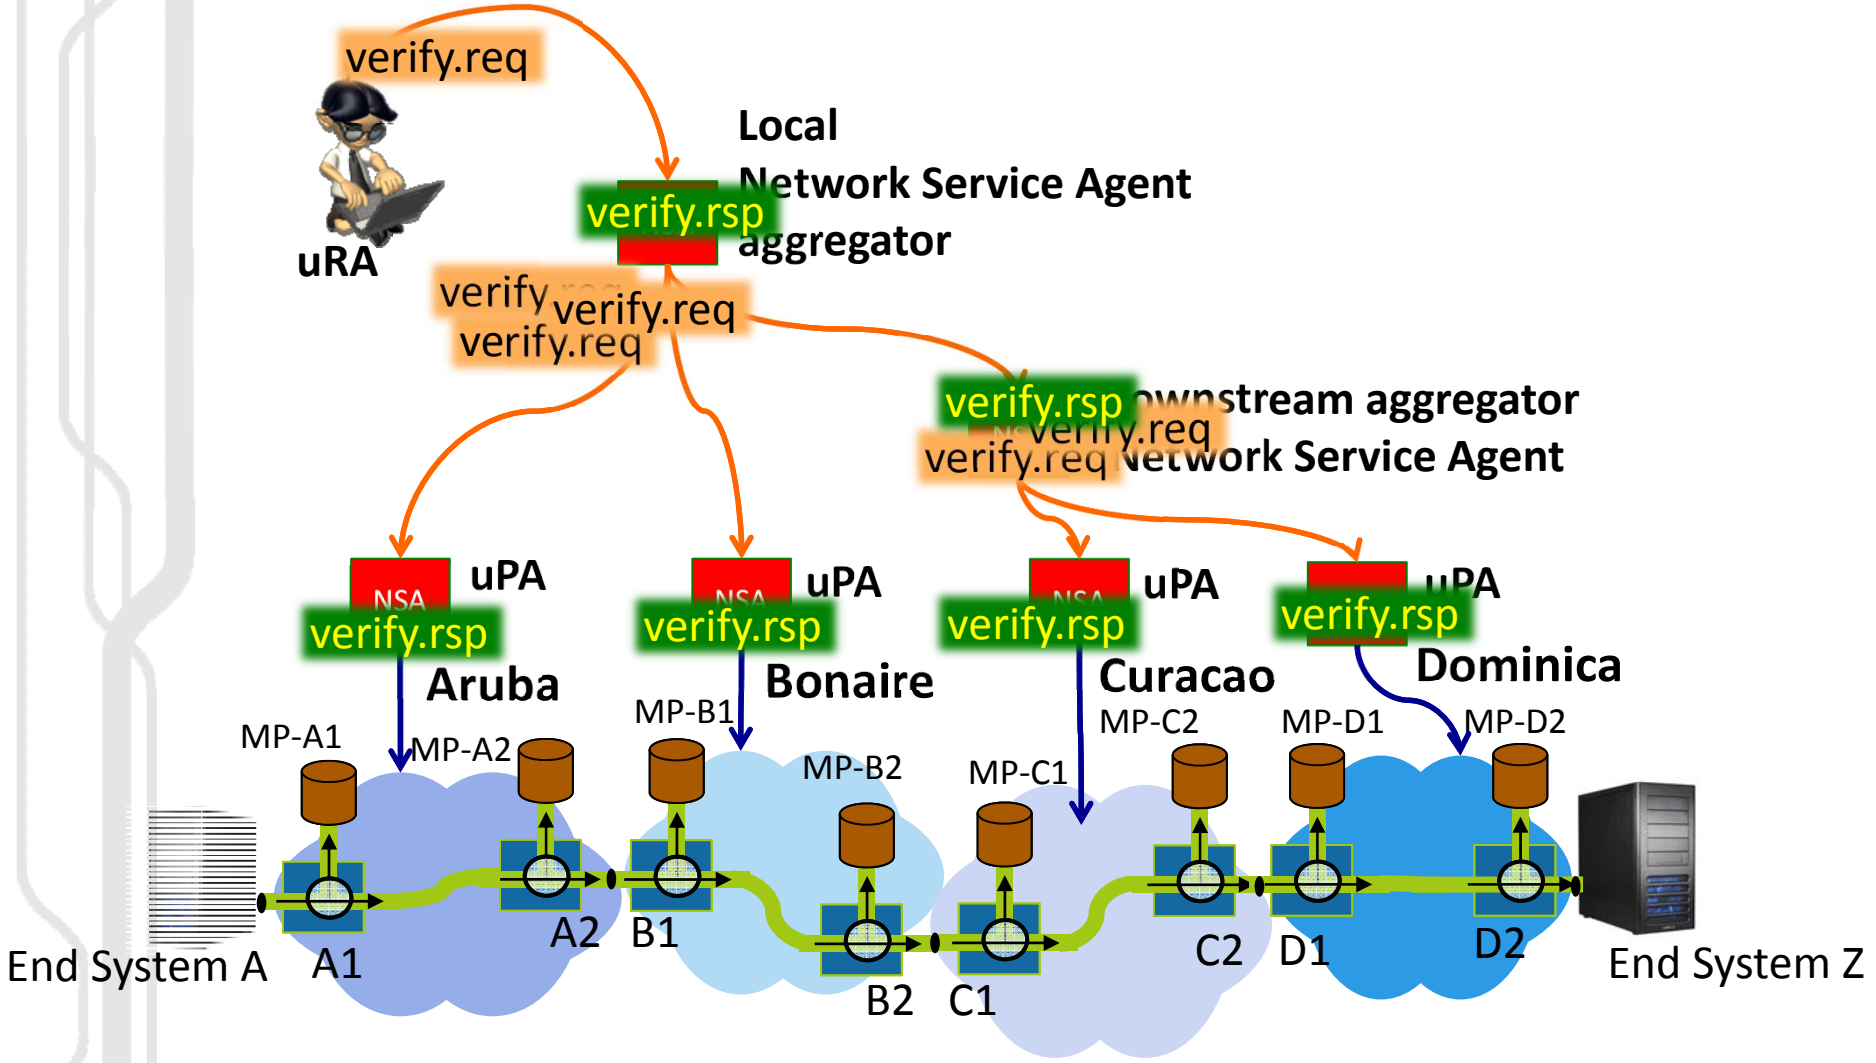

NORDUnet
Nordic infrastructure for Research & Education

Summary

- PV for performance guaranteed services involves understanding the following:
    - A) What the service purports to be able to do
    - B) What the service instance has been guaranteed
    - C) What we can effectively test
  - New technologies:
    - Advanced service architectures
    - Advanced and well engineered network/data plane architectures
    - Advanced high resolution timing and interface technologies.
  - Use existing tools and packages where they can fit into the picture
- Use the GLIF PVA Task Force to define a “stake in the sand” from which to start.

- The End