DetNet   J. Farkas
Internet-Draft   B. Varga
Intended status: Standards Track   Ericsson
Expires: September 13, 2017   R. Cummings
                                                  National Instruments
                                                        March 12, 2017

                DetNet Flow Information Model Based on TSN
                draft-farkas-detnet-flow-information-model-00

Abstract

   This document describes flow information model for Deterministic
   Networking (DetNet). The DetNet service is provided either for a
   Layer 3 or a Layer 2 flow. This document provides DetNet flow
   information model both for Layer 3 and Layer 2 flows in an integrated
   fashion.

Status of This Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at http://datatracker.ietf.org/drafts/current/.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   This Internet-Draft will expire on September 13, 2017.

Copyright Notice

Table of Contents

1. Introduction

   A Deterministic Networking (DetNet) service provides a capability to
   carry a unicast or a multicast data flow for an application with
   constrained requirements on network performance, e.g., low packet
   loss rate and/or latency. The DetNet service is provided either for
   a Layer 3 (L3) flow or a Layer 2 (L2) flow by an IP/MPLS network,
   see, e.g., [I-D.ietf-detnet-dp-alt]. Similarly, Time-Sensitive
   Networking (TSN) [IEEE8021TSN]) can be used for L2 flows in a bridged
   network. DetNet and TSN have common architecture as expressed in
   [IETFDetNet] and [I-D.ietf-detnet-architecture]. DetNet service can
   be leveraged both by L3 and L2 flows, i.e., by DetNet L3 flows and

DetNet L2 flows. Therefore, the DetNet flow information model
provided by this document covers both DetNet L3 flows and DetNet L2
flows in an integrated fashion. Thus, the DetNet flow information
model is based on [I-D.ietf-detnet-architecture] and on the data
model specified by [IEEE8021Qcc]. Furthermore, the DetNet flow
information model relies on the flow identification possibilities
described in [IEEE8021CB], which is used by [IEEE8021Qcc] as well.
In addition to TSN data model, [IEEE8021Qcc] also specifies
configuration of TSN features (e.g., traffic scheduling specified by
[IEEE8021Qbv]). Due to the common architecture and flow model,
configuration features can be leveraged in certain deployment
scenarios, e.g., when the network that provides the DetNet service
includes both L3 and L2 network segments.

Based on the DetNet architecture [I-D.ietf-detnet-architecture] (see
Section 4), this document (this revision) only considers the
Centralized Network / Distributed User Model out of the models
specified by [IEEE8021Qcc]. That is, there is a User-Network
Interface (UNI) between an end system and a network. Furthermore,
there is a central entity for the control of the network. For
instance, the central entity implements a Path Computation Element
(PCE) for the calculation and establishment of paths needed for
packet replication and elimination, if any.

[[NOTE (to be removed from a future revision): The Goals and Non
goals subsections are only for revision 00, they are to be removed
from future revisions of this draft.]]

1.1. Goals

As it is expressed in the Charter [IETFDetNet], the DetNet WG
collaborates with IEEE 802.1 TSN in order to define a common
architecture for both Layer 2 and Layer 3, which is beneficial for
various reasons, e.g., in order to simplify implementations. The
flow information model should be also common along those lines. As
the TSN flow information/data model specified by [IEEE8021Qcc] is
mature, the DetNet flow information model described in this document
is based on [IEEE8021Qcc], which is an amendment to [IEEE8021Q].

The Centralized Network / Distributed User Model of [IEEE8021Qcc] is
used in this revision as a start of the work. Further models can be
also useful for DetNet, e.g., the Fully Centralized Model for the
Industrial M2M use case [I-D.ietf-detnet-use-cases].

This document intends to specify flow information model only.

Revision 00 is just a start; it is not complete. As this revision
heavily relies on [IEEE8021Qcc], the need for further DetNet specific
aspects is to be reviewed and missing pieces are to be added.

1.2. Non Goals

This document (this revision) does not intend to specify either flow
data model or DetNet configuration. From these aspects, the goals of
this document differ from the goals of [IEEE8021Qcc], which also
specifies data model and configuration of certain TSN features.

2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in [RFC2119].

The lowercase forms with an initial capital "Must", "Must Not",
"Shall", "Shall Not", "Should", "Should Not", "May", and "Optional"
in this document are to be interpreted in the sense defined in
[RFC2119], but are used where the normative behavior is defined in
documents published by SDOs other than the IETF.

3. Terminology and Definitions

This document uses the terminology established in Section 2 of the
DetNet architecture document [I-D.ietf-detnet-architecture]. The
DetNet <=> TSN dictionary of [I-D.ietf-detnet-architecture] is used
to perform translation from [IEEE8021Qcc] to this document.
Additional terms used in this document:

DetNet L3 Flow: Layer 3 (L3) flow leveraging DetNet service.

DetNet L2 Flow: Layer 2 (L2) flow leveraging DetNet service.

4. Naming Conventions

The following naming conventions were used for naming information
model components in this document. It is recommended that extensions
of the model use the same conventions.

o Names SHOULD be descriptive.

o Names MUST start with uppercase letters.

o Composed names MUST use capital letters for the first letter of
   each component. All other letters are lowercase, even for
   acronyms. Exceptions are made for acronyms containing a mixture

of lowercase and capital letters, such as IPv6. Examples are
SourceMacAddress and DestinationIPv6Address.

5. End System

   Deterministic service is required by time/loss sensitive
   application(s) running on an end system during communication with its
   peer(s). Such a data exchange has various requirements on delay and/
   or loss parameters.

   The DetNet architecture [I-D.ietf-detnet-architecture] distinguishes
   two kinds of end systems: Source and Destination. The same
   distinction is applied for the DetNet flow information model. In
   addition to the end systems interested in a flow, the status
   information of the flow is also important. Therefore, the DetNet
   flow information model relies on three high level groups:

   o Source: an end system capable of sourcing a DetNet flow. The
      Source information group includes elements that specify the Source
      for a single flow. This information group is applied from the
      user to the network.

   o Destination: an end system that is a destination of a DetNet flow.
      The Destination information group includes elements that specify
      the Destination for a single flow. This information group is
      applied from the user to the network.

   o Status: the status of a DetNet flow. The status information group
      includes elements that specify the status of the flow in the
      network. This information group is applied from the network to
      the user. This information group informs the user whether or not
      the flow is ready for use.

   There are two operations for each flow with respect to a Source or a
   Destination:

   o Join: Source/Destination request to join the flow.

   o Leave: Source/Destination request to leave the flow.

   [[NOTE (to be removed from a future revision): Adding Modify
   operation can be considered to address cases when a flow is slightly
   changed, e.g., only MaxPacketSize (Section 6.2) has been changed.]]

   As the DetNet UNI can provide both L3 and L2 services, end systems
   may not need to implement the L3 <=> L2 Transfer Function specified
   by [IEEE8021CB] (see, e.g., subclause 6.3; see also subclause 46.1 in
   [IEEE8021Qcc]). An edge node may implement a function similar to the

Transfer Function, see, e.g., the Svc Proxy in Figure 1 in
[I-D.ietf-detnet-dp-alt].

6. Flow

The flows leveraging DetNet service can be unicast or multicast data
flows for an application with constrained requirements on network
performance, e.g., low packet loss rate and/or latency. Therefore,
they can require different connectivity types: point-to-point (p2p)
or point-to-multipoint (p2mp). The p2mp connectivity is created by a
transport layer function (e.g., p2mp LSP) [I-D.ietf-detnet-dp-alt].
(Note that mp2mp connectivity is a superposition of p2mp
connections.)

Many flows using DetNet service are periodic with fix packet size
(i.e., Constant Bit Rate (CBR) flows), or periodic with variable
packet size.

Delay and loss parameters are correlated because the effect of late
delivery can result data loss for an application. However, not all
applications require hard limits on both parameters (delay and loss).
For example, some real-time applications allow graceful degradation
if loss happens (e.g., sample-based processing, media distribution).
Some others may require high-bandwidth connections that make the
usage of techniques like packet replication economically challenging
or even impossible. Some applications may not tolerate loss, but are
not delay sensitive (e.g., bufferless sensors). Time/loss sensitive
applications may have somewhat special requirements especially for
loss (e.g., no loss in two consecutive communication cycles; very low
outage time, etc.).

Flows have the following attributes:

a. DataFlowSpecification (Section 6.1)

b. TrafficSpecification (Section 6.2)

c. FlowRank (Section 6.3)

Flow attributes are described in the following sections.

6.1. Identification and Specification of Flows

Identification options for TSN flows are specified by [IEEE8021CB],
which also includes IP flow identification, see, e.g., Table 6-1 in
Clause 6. Therefore, the flow identification specified by
[IEEE8021CB] is also applicable to DetNet flows.

[[NOTE (to be removed from a future revision): Extensions to the
options specified by [IEEE8021CB] can be discussed.]]

DataFlowSpecification specifies DetNet flows as follows; see
Section 6.1.1 for DetNet L3 flows and Section 6.1.2 for DetNet L2
flows.

6.1.1. DetNet L3 Flow Identification and Specification

DetNet L3 flows can be identified and specified by the following
attributes:

a. SourceIpAddress

b. DestinationIpAddress

c. Dscp

d. Protocol

e. SourcePort

f. DestinationPort

g. MplsLabel

6.1.2. DetNet L2 Flow Identification and Specification

DetNet L2 flows can be identified and specified by the following
attributes:

a. DestinationMacAddress

b. SourceMacAddress

c. Pcp

d. VlanId

[[NOTE (to be removed from a future revision): The Multiple Stream
Registration Protocol (MSRP) [IEEE8021Q] uses StreamID to match
Talker registrations with their corresponding Listener registrations,
i.e., to identify Streams (L2 TSN flows). The StreamID includes the
following subcomponents:

o A 48-bit MAC Address associated with the Talker sourcing the
   stream to the bridged network.

   o A 16-bit unsigned integer value, Unique ID, used to distinguish
      among multiple streams sourced by the same Talker.

   ]]

6.2. Traffic Specification

   TrafficSpecification specifies how the Source transmits packets for
   the flow. This is effectively the promise/request of the Source to
   the network. The network uses this traffic specification to allocate
   resources and adjust queue parameters in network nodes.

   TrafficSpecification has the following attributes:

   a. Interval: the period of time in which the traffic specification
       cannot be exceeded.

   b. MaxPacketsPerInterval: the maximum number of packets that the
       Source will transmit in one Interval.

   c. MaxPayloadSize: the maximum payload size that the Source will
       transmit.

6.3. Flow Rank

   FlowRank provides the rank of this flow relative to others flows in
   the network. This rank is used to determine success/failure of flow
   establishment. Rank (boolean) is used by the network to decide which
   flows can and cannot exist when network resources reach their limit.
   Rank is used to help to determine which flows can be dropped (i.e.,
   removed from node configuration) if a port of a node becomes
   oversubscribed (e.g., due to network reconfiguration). The false
   value is more important than the true value (i.e., flows with true
   are dropped first).

7. Source

   The Source object specifies:

   o The behavior of the Source for the flow (how/when the Source
      transmits).

   o The requirements of the Source from the network.

   o The capabilities of the interface(s) of the Source.

   The Source object includes the following attributes:

   a. DataFlowSpecification (Section 6.1)

   b. TrafficSpecification (Section 6.2)

   c. FlowRank (Section 6.3)

   d. EndSystemInterfaces (Section 9.1)

   e. InterfaceCapabilities (Section 9.2)

   f. UserToNetworkRequirements (Section 9.3)

   For the join operation, the DataFlowSpecification, FlowRank,
   EndSystemInterfaces, and TrafficSpecification SHALL be included
   within the Source. For the join operation, the
   UserToNetworkRequirements and InterfaceCapabilities groups MAY be
   included within the Source.

   For the leave operation, the DataFlowSpecification and
   EndSystemInterfaces SHALL be included within the Source.

8. Destination

   The Destination object includes the following attributes:

   a. DataFlowSpecification (Section 6.1)

   b. EndSystemInterfaces (Section 9.1)

   c. InterfaceCapabilities (Section 9.2)

   d. UserToNetworkRequirements (Section 9.3)

   For the join operation, the DataFlowSpecification and
   EndSystemInterfaces SHALL be included within the Destination. For
   the join operation, the UserToNetworkRequirements and
   InterfaceCapabilities groups MAY be included within the Destination.

   For the leave operation, the DataFlowSpecification and
   EndSystemInterfaces SHALL be included within the Destination.

   [[NOTE (to be removed from a future revision): Should we add
   DestinationRank? It could distinguish the importance of Destinations
   if the flow cannot be provided for all Destinations.]]

9. Common Attributes of Source and Destination

   Source and Destination end systems have the following common
   attributes in addition to DataFlowSpecification (Section 6.1).

9.1. End System Interfaces

   EndSystemInterfaces is a list of identifiers, one for each physical
   interface (port) in the end system acting as a Source or Destination.
   An interface is identified by an IP or a MAC address.

   [[NOTE (to be removed from a future revision): Sub-Interfaces to be
   added, e.g., based on IfIndex.]]

9.2. Interface Capabilities

   InterfaceCapabilities specifies the network capabilities of all
   interfaces (ports) contained in the EndSystemInterfaces object
   (Section 9.1). These capabilities may be configured via the
   InterfaceConfiguration object (Section 10.2) of the Status object
   (Section 10).

   Note that an end system may have multiple interfaces with different
   network capabilities. In this case, each interface should be
   specified in a distinct top-level Source or Destination object (i.e.,
   one entry in EndSystemInterfaces (Section 9.1)). Use of multiple
   entries in EndSystemInterfaces is intended for network capabilities
   that span multiple interfaces (e.g., packet replication and
   elimination)."；.

   [[NOTE (to be removed from a future revision): InterfaceCapabilities
   attributes are to be defined. For information, [IEEE8021Qcc]
   specifies the following attributes:

   a. VlanTagCapable (Customer VLAN Tag capable)

   b. CB-Capable (frame replication and elimination capable)

   c. CB-StreamIdenTypeList (a list of the optional Stream
       Identification types supported by the interface as specified in
       [IEEE8021CB].)

   d. CB-SequenceTypeList (a list of the optional Sequence Encode/
       Decode types supported by the interface as specified in
       [IEEE8021CB].)

   ]]

9.3. User to Network Requirements

   UserToNetworkRequirements specifies user requirements for the flow,
   such as latency and reliability.

   The UserToNetworkRequirements object includes the following
   attributes:

   a. NumReplicationTrees

   b. MaxLatency

   NumReplicationTrees specifies the number of maximally disjoint trees
   that the network should configure to provide packet replication and
   elimination for the flow. NumReplicationTrees is provided by the
   Source only. Destinations SHALL set this element to one. Value zero
   and one indicate no packet replication and elimination for the flow.
   When NumReplicationTrees is greater than one, packet replication and
   elimination is to be used for the flow. If the Source sets this
   element to greater than one, and packet replication and elimination
   is not possible in the network (e.g., no disjoint paths, or the nodes
   do not support packet replication and elimination), then the
   FailureCode of the Status object is non-zero (Section 10.1).

   MaxLatency is the maximum latency from Source to Destination(s) for a
   single packet of the flow. MaxLatency is specified as an integer
   number of nanoseconds. When this requirement is specified by the
   Source, it must be satisfied for all Destinations. When this
   requirement is specified by a Destination, it must be satisfied for
   that particular Destination only. If the UserToNetworkRequirements
   group is not provided within the Source or Destination object, then
   value zero SHALL be used for this element. Value zero represents a
   special use for the maximum latency requirement. Value zero locks-
   down the initial latency that the network provides in the
   AccumulatedLatency parameter of the Status object (Section 10) after
   the successful configuration of the flow, such that any subsequent
   increase in the latency beyond that initial value causes the flow to
   fail.

   [[NOTE-1 (to be removed from a future revision): Should we add a
   parameter to specify the maximum packet loss rate that can be
   tolerated for the flow?]]

   [[NOTE-2 (to be removed from a future revision): TrafficSpecification
   (Section 6.2) specifies the Peak Information Rate (PIR) of the flow,
   which is a kind of user requirement to the network. Should we add
   Committed Information Rate (CIR), i.e., the minimum rate the user
   requests to be guaranteed for the flow by the network?]]

10. Status

   The Status object is provided by the network each Source and
   Destination of the flow. The Status object provides the status of
   the flow with respect to the establishment of the flow by the
   network. The Status object is delivered via the corresponding UNI to
   each Source and Destination end system of the flow. The Status is
   distinct for each Source or Destination because the
   AccumulatedLatency and InterfaceConfiguration objects are distinct,
   see below.

   The Status object SHALL include the attributes a), b), c); and MAY
   include attributes d), e):

   a. DataFlowSpecification (Section 6.1)

   b. StatusInfo (Section 10.1)

   c. AccumulatedLatency (this section below)

   d. InterfaceConfiguration (Section 10.2)

   e. FailedInterfaces (Section 10.3)

   DataFlowSpecification identifies the flow for which status is
   provided. DataFlowSpecification is described in (Section 6.1) If the
   Status object is provided without a Source or Destination object in a
   protocol message via a UNI, then the DataFlowSpecification object
   SHALL be included within the Status object for both join and leave
   operations. If the Status object immediately follows a Source or
   Destination object in the protocol message, then the
   DataFlowSpecification object is obtained from the Source/Destination
   object, and therefore DataFlowSpecification is not required within
   the Status object.

   AccumulatedLatency provides the worst-case latency that a single
   packet of the flow can encounter along its current path(s) in the
   network. When provided to a Source, AccumulatedLatency is the worst-
   case latency for all Destinations (worst path). AccumulatedLatency
   is specified as an integer number of nanoseconds. Latency is
   measured using the time at which the data frame's message timestamp
   point passes the reference plane marking the boundary between the
   network media and PHY. The message timestamp point is specified by
   IEEE Std 802.1AS [IEEE8021AS] for various media. For a successful
   Status, the network returns a value less than or equal to the
   MaxLatency of the UserToNetworkRequirements (Section 9.3). If the
   NumReplicationTrees of the UserToNetworkRequirements (Section 9.3) is
   one, then the AccumlatedLatency SHALL provide the worst latency for

the current path from the Source to each Destination. If the path is changed (e.g., due to rerouting), then the AccumulatedLatency changes accordingly. If the NumReplicationTrees of the UserToNetworkRequirements (Section 9.3) is greater than one, AccumlatedLatency SHALL provide the worst latency for all paths in use from the Source to each Destination.

10.1. Status Info

   StatusInfo provides information regarding the status of a flow's configuration in the network.

   The StatusInfo object MAY include the following attributes:

   a. SourceStatus is an enumeration for the status of the flow's Source:

       * None: no Source

       * Ready: Source is ready

       * Failed: Source failed

   b. DestinationStatus is an enumeration for the status of the flow's Destinations:

       * None: no Destination

       * Ready: all Destinations are ready

       * PartialFailed: One or more Destinations ready, and one or more Listeners failed. The flow can be used tf the Source is Ready.

       * Failed: All Destinations failed.

   c. FailureCode: A non-zero code that specifies the problem if the flow encounters a failure (e.g., packet replication and elimination is requested but not possible, or SourceStatus is Failed, or DestinationStatus is Failed, or DestinationStatus is PartialFailed).

   [[NOTE (to be removed from a future revision): FailureCodes to be defined for DetNet. Table 46-1 of [IEEE8021Qcc] describes TSN failure codes.]]

10.2. Interface Configuration

   InterfaceConfiguration provides configuration of interfaces in the
   Source/Destination. This configuration assists the network in
   meeting the requirements of the flow. The InterfaceConfiguration
   object is according to the capabilities of the interface.
   InterfaceConfiguration can be distinct for each Source or Destination
   of each flow. If the InterfaceConfiguration object is not provided
   within the Status object, then the network SHALL assume zero elements
   as the default (no interface configuration).

   The InterfaceConfiguration object MAY include one or more the
   following attributes:

   a. MAC or IP Address to identify the interface

   b. DataFlowSpecification (Section 6.1)

10.3. Failed Interfaces

   FailedInterfaces provides a list of one or more physical interfaces
   (ports) in the failed node when a failure occurs in network
   configuration (i.e., non-zero FailureCode in StatusInfo object
   (Section 10.1)).

   The InterfaceConfiguration object includes the following attributes:

   a. MAC or IP Address to identify the interface

   b. InterfaceName

   InterfaceName is the name of the interface (port) within the node.
   This interface name SHALL be persistent, and unique within the node.

11. Summary

   This document describes DetNet flow information model both for DetNet
   L3 flows and DetNet L2 flows based on the TSN data model specified by
   [IEEE8021Qcc]. This revision of the document is just to start the
   discussions; further work is needed.

12. IANA Considerations

   N/A.

13. Security Considerations

   N/A.

14. References

14.1. Normative References

   [I-D.ietf-detnet-architecture]
             Finn, N. and P. Thubert, "Deterministic Networking
             Architecture", draft-ietf-detnet-architecture-00 (work in
             progress), September 2016.

   [I-D.ietf-detnet-dp-alt]
             Korhonen, J., Farkas, J., Mirsky, G., Thubert, P.,
             Zhuangyan, Z., and L. Berger, "DetNet Data Plane Protocol
             and Solution Alternatives", draft-ietf-detnet-dp-alt-00
             (work in progress), October 2016.

   [I-D.ietf-detnet-use-cases]
             Grossman, E., Gunther, C., Thubert, P., Wetterwald, P.,
             Raymond, J., Korhonen, J., Kaneko, Y., Das, S., Zha, Y.,
             Varga, B., Farkas, J., Goetz, F., and J. Schmitt,
             "Deterministic Networking Use Cases", draft-ietf-detnet-
             use-cases-01 (work in progress), February 2016.

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
             Requirement Levels", BCP 14, RFC 2119,
             DOI 10.17487/RFC2119, March 1997,
             <http://www.rfc-editor.org/info/rfc2119>.

14.2. Informative References

   [IEEE8021AS]
             IEEE 802.1, "IEEE 802.1AS-2011: IEEE Standard for Local
             and metropolitan area networks - Timing and
             Synchronization for Time-Sensitive Applications in Bridged
             Local Area Networks", 2011,
             <http://standards.ieee.org/getieee802/
             download/802.1AS-2011.pdf>.

   [IEEE8021CB]
             IEEE 802.1, "IEEE P802.1CB: IEEE Draft Standard for Local
             and metropolitan area networks - Frame Replication and
             Elimination for Reliability", 2017,
             <http://www.ieee802.org/1/pages/802.1cb.html>.

[IEEE8021Q]
          IEEE 802.1, "IEEE 802.1Q-2014: IEEE Standard for Local and
          metropolitan area networks - Bridges and Bridged
          Networks", 2014, <http://standards.ieee.org/getieee802/
          download/802-1Q-2014.pdf>.

[IEEE8021Qbv]
          IEEE 802.1, "IEEE 802.1Qbv-2015: IEEE Standard for Local
          and metropolitan area networks - Bridges and Bridged
          Networks -- Amendment 25: Enhancements for Scheduled
          Traffic", 2015, <https://standards.ieee.org/findstds/
          standard/802.1Qbv-2015.html>.

[IEEE8021Qcc]
          IEEE 802.1, "IEEE P802.1Qcc-2015: IEEE Draft Standard for
          Local and metropolitan area networks - Bridges and Bridged
          Networks -- Amendment: Stream Reservation Protocol (SRP)
          Enhancements and Performance Improvements", 2017,
          <http://www.ieee802.org/1/pages/802.1cc.html>.

[IEEE8021TSN]
          IEEE 802.1, "IEEE 802.1 Time-Sensitive Networking (TSN)
          Task Group", <http://www.ieee802.org/1/pages/tsn.html>.

[IETFDetNet]
          IETF, "IETF Deterministic Networking (DetNet) Working
          Group", <https://datatracker.ietf.org/wg/detnet/charter/>.

Authors' Addresses

   Janos Farkas
   Ericsson
   Konyves Kalman krt. 11/B
   Budapest 1097
   Hungary

   Email: janos.farkas@ericsson.com

   Balazs Varga
   Ericsson
   Konyves Kalman krt. 11/B
   Budapest 1097
   Hungary

   Email: balazs.a.varga@ericsson.com

Rodney Cummings
National Instruments
11500 N. Mopac Expwy
Bldg. C
Austin, TX 78759-3504
USA

Email: rodney.cummings@ni.com