NSIS interim meeting Mon10-Wed12

(Monday February 10)

Attendants:

Marcus Brunner, Ruediger Geib, Hans Lippitsch, Hannes Tschofenig, Henning Schulzrinne, Robert Hancock, Kwok Ho Chan, Ping Pan, Joachim Hillbrandt, Janne Rinne, Scott Bradner

Minute Taker:

Hannes Tschofenig; Sven van den Bosch

Agenda:

WG update
skipped

Agenda Bashing
skipped

Introduction of the participants

see participants list

Markus Brunner: Requirements update

John created an open issues list for NSIS:
http://www-nrc.nokia.com/sua/nsis/nsis-issues.htm
Markus went through that list and addresses each open item

Markus: Requirements draft is intentionally kept abstract.

Issue 1:

ACTION: Remove text 5.1.5 and check exclusion section
REASON: out-of-scope

Issue 2:

DISCUSSION:
Henning: The requirements and the framework document should not be conflicting. Referring to the framework is possible.
Term is never mentioned in the document.

DISCUSSION:
There are no requirements about sender- or receiver initiated signaling.

5.6.2 (Flexibility in the placement of the NSIS Initiator) is related to this.
Two issues:
- Directionality
- Placement

Henning: State maintenance and setup
Signaling protocol does
- control protocol which visits the path along the data path (or follows it roughly) ~ where is goes
- establishes/modify/remove state at some nodes (in most cases - it could also only query)
  * state necessary for the layer split (route messages from a to b)
  * state for the application state (e.g. midcom, QoS, etc.)

state definition: sticks around after the packet is gone

active networking is such an example (no resource in the traditional sense)

should there be a generic protocol requirements

robert: how far do we want the requirements draft to go? making the requirements document independently of any type of application is a very hard job.
proposal: requirements draft is generic (introduction) but the examples are qos driven

ACTION: remove sender / receiver initiated signaling definitions /
reference 5.6.2 what is intended.

It should be tried to a definition of state.

Issue 3: 5.2.2
DISCUSSION:
Henning: It is only a matter of how closely on the data path (not on-path / off-path)

Testable definition:
a) same "virtual" interfaces
b) same ASes
(these are the only terms defined for the internet routing architecture)

we could use framework definition in section 3.1.1 / path couples - path decoupled

RSVP-TE uses RSVP in the reverse order (signaling establishes routing)

Robert: Look into the framework draft and use definition there.
must support path-coupled, should not exclude path-decoupled.

there is no term on-path (only path-coupled)!

ACTION:
The agreement is to use the framework terminology and definitions of
'path-coupled' and 'path-decoupled'. The requirement then becomes. 'The
protocol MUST support path-coupled and SHOULD NOT exclude path-decoupled
signaling'.

Issue 4: 5.3.4

DISCUSSION:
Henning: What is a request for service? More an application issue? This requirement is very QoS biased.
Message returned has two objects:
a) Confirms the establishment of resource reservation
b) Reliability requirement : NSIS message has been transmitted somewhere (and we got a return message back)

There is no clear understanding of the requirement.

Henning: Confirmation of service execution / Configuration of reliability
Hence this requirement can be split into two separate requirements.

Robert: This requirement cannot be seen in isolation (5.10.1).
Reliability does not need to be addressed.

It should be possible to obtain indication of success/failure.
Requirement for a specific service layer (since there it only means something).

It should be possible to request a response message.

Who has to answer the question.

ACTION: Shuffle things around.
There are two aspects to the proposed text:
a. confirmation of delivery
b. confirmation of service-related action (e.g. installing state)
The first one is covered by 5.10.1. It is proposed to delete the text proposed for this issue (5.3.4) but insert a requirement covering the second aspects after 5.10.2. Ruediger will provide text.

Issues 5:

Meaning is somewhat vague.

Marcus:
- RSVP bundling type of requirement
- Not about aggregation.

Paragraph is QoS reservation centric.
Make it more generic? change reservation to flow?

ACTION: replace "MUST NOT know" with "need not to know"
Remove second half of the last sentence.

Issue 6:
ACTION: no conclusion - skipped

Issue 7: UMTS access session

ACTION: already done - agreed.

Issue 8: Definition of RMF

ACTION: agreed

Issue 9: Provisioning text

ACTION: remove text about provisioning

Issue 10: QoS technology

ACTION: remove

Issue 11: Clarification on NSIS usage

ACTION: send email to requestor of item and ask where to place the text
add middlebox communication as an example of a further NSIS signaling application

Issue 12: Requirement on routing

NSIS does not interfere with routing.
Henning: determination of next data node selection is not done by NSIS
(forwarding is done by someone else)
(next hop of the data packet)

Background: NSIS is not a QoS routing protocol; relationship to load balancing
Relationship to 5.9.6 (non-traditional routing)

ACTION:
Replace text with "NSIS assumes layer 3 routing and the determination of next data node selection is not done by NSIS".

Issue 13: Clarification

ACTION: accept

Issue 14: 5.2.2

ACTION: nothing. the text has changed already.

some minor editorial issues

Issue 16:

ACTION: accepted

Issues 17:

ACTION: Remove paragraph

Issues 18:

ACTION: remove

Issue 19:

ACTION: remove point 7 of protection of non-signaling messages
A requirement for encapsulation is covered with the flow identification requirement (5.9.1)

Issue 20: 5.1.1

ACTION: remove

Issue 21:
5.3.3 (NSIS SHOULD allow for sending notifications upstream)

ACTION: don't change

Issue 22:
ACTION: change accepted

Issue 23:
ACTION: editorial changes

Issue 24: 5.1.7
DISCUSSION:
Problem caused through terminology of two-layer split in the framework document.
framework:
a) signaling application (qos midcom)
b) user application (for user application triggering something).
Henning: meaning is confusing.
Opaque application information MAY get transported in the signaling message, without being handled in the network.
Two issues:
a) only end-to-end object (why is it carried in NSIS?)
b) only interpreted by some entities (the nsis protocol should be able to pass around objects which are interpreted only by some nodes)

ACTION: remove 5.1.7
add paragraph from above (see b): something like: "NSIS should be able to carry opaque objects. These are objects that are only interpreted at some NSIS-capable nodes."

Issue 25:
DISCUSSION:
Should the requirements document reference the framework document./
Henning: Informative reference is not going to block.

ACTION: add the reference as an informative reference

Generic Requirements Discussion

Henning: comment the requirements draft.
Indicate that some sections are generic whereas others are specific to an application
Add a separate section to cover specific issues of midcom, qos, etc.
The terminology "reservation" used in the requirements document has a QoS bias. The requirement should talk about state information.

Henning:
Question: Identify requirements that are application signaling specific

Robert:
For an outsider the document is not too easy to understand.
There should be a uniform label which is assigned to each section.

People fear that this activity stalls the progress of the document.

Henning:
Possible options:
a) labeling (applicability statement)
b) grouping it to sections

What is the requirements document:
- working document to capture discussion (people might argue with it as a design decision)
- outsiders to understand

Possible outcomes of a labeling process:
- for some things it might be very clear
- some things we don't understand (so vage to be useless)

Problems:
- Vague document -> danger of wrong interpretation
- if we are not able as a working group then the req. needs to be dropped or re-defined.

Henning:
=> Homework exercise: go through the document and label different sections.

2. Security threats for NSIS - update (Hannes)

Hannes goes through his presentation.

Issues:
Agreement on the minimum security requirements (authentication, integrity and replay protection) on a peer-to-peer basis

Should key management be addressed by NSIS protocols?
-> yes

Discussion on whether there should be the assumption that the next NSIS peer is known.
It cannot be assumed in all circumstances. There are some NSIS application (e.g. middlebox communication) and some environments where such an assumption cannot be made.

Discussion on end-to-end security
-> End-to-end security is an NSLP concept. It has to be used only if the semantic of the object requires it. In general NSIS should not be used to solely exchange information end-to-end. As a criteria one could imagine that the protected objects is not visible to the intermediate nodes. If this object is not relevant for NSIS intermediate nodes then it should be exchanged with an end-to-end protocol only (such as SIP, http, etc.)

Discussion on mobility: it was agreed that mobility is an open issue that needs to be tackled in the framework context first. Two questions need to be addressed:
a. is the session identifier useful as mechanism (for security, ...)
b. if security-relevant, what do we need to do to protect it (ill-defined for now: who sets it, ...)

Implications of Trust Relationships for NSIS Signaling (Hannes)

Presentation addresses middlebox signaling (firewall; NAT)

Goal: figure out difference between QoS and middlebox signaling (related to TIST)

Hannes goes through the presentation

What happens next?
Hannes: we talked about extending the document (with Cedric Aoun) and making it more generic

Is the goal to come up with requirements specific to for a middlebox communication NSLP application?
Cedric: we already have requirements and framework in midcom. The proposal is to have an NSLP for midcom, and look at NTLP issues from a midcom point-of-view.

Is the CASP document contradicting/duplicating/different from work done in the midcom working group?

-> It is a good start but need to be more generic and add some things to it.
Remark: The midcom wg tries to solve the problem with a different focus and different working assumptions (path decoupled vs. path decoupled signaling)

Henning: Since the work in the midcom group is different we should not use the term midcom but rather 'middlebox configuration' to avoid confusion
Remark: Would like to capture specificity of firewall/NAT to reduce scope.

We are going to look around for other applications on wednesday morning

NSIS AAA issues (Hannes)

Hannes goes through the presentation
Q: Do you assume per-flow accounting between domains? That will certainly not be the case.
Remark: Parkway is logically not that different from roaming user where all nodes can be visited networks.
Remark: Transit providers do not have the AAA infrastructure (per-flow, roaming, ...) required for the Parkway model
Q: Why is Parkway model even considered?
A: Because VoIP uses clearing house and because at least it allows for fixed local prices
Q: On which layer does it apply?

A: The upper layer
Remark: The impact on NTLP is the addition of another dimension to data sender/receiver, signaling initiator/responder, now another one is added (initiator/receiver charged)
Q: Have you looked at DoS attacks for security?
A: This is part of the threat analysis
Remark: The AAA consideration are important from trust relations point of view
Framework presentation (Robert)

Request for mobility and signaling application space (i.e. other non-qos applications) release for next ietf (sfo)

Key Issue 1:

- We don't have to put all generic protocols in the ntlp.
- Capturing the split with an api?

Key Issues 2:

- not much discussion on the mobility issues (no conclusions - what and how todo)
interaction with seamoby activities is unclear

Key Issues 3:

Shift toward multi-application protocol.
What needs to be done to prevent chaos?

Issue 1:

Sender-/Receiver definition is unclear
NI/NF/NR become NSLP concepts

Issue 2:

Who is allowed what todo for a reservation?
Can the network change something?
Is only the initiator allowed to change something?
It is about authorization!

Proxy:

Node along the path that "return" signaling message (case b) or node which acts on behalf of the sender (case a).

The requirements/framework document already allows (a). Case (b) raises more discussion.

The term of proxy is somewhat a confusion attractor. The term proxy is also (by some people) used as a synonym for a bandwidth broker.

Issue 3:

Should entities only related to their neighboring peers?
Should the security go over more than a single hop (notification handling - though various hops or to the remote end,

Properties of the session identifier:
- constant over the lifetime of a session
- selected "magically" by the initiator

Functionality of the NTLP is a layer split discussion.

Issue 4:

Addressing (peer-to-peer or end-to-end)
Certain messages are end-to-end (e.g. discovery messages)

Issue 5:

Flow Identification

What parameters to use?
Current assumption: something minimal is at the NTLP
HoA for MIP support is currently ruled out.
Problems with policy based forwarding will be explained.

Flow Identification at the NLSP requires that routing is done at the application layer (also NAT handling)

Relationship with session identifier pointed out.

Issue 6:

Mobility issues: many things unclear (micro-/macro-mobility), which optimizations are possible/useful?

Henning: Minimal requirement (if a flow identifier is changed then existing state needs to be recognized).

Relationship between the flow identifier and the mobility handling

Issue 7:
see slides
Issue 8:

Path-decoupledness

Can be discussed in the two-layer discussion

RSVP Transport Issues (Ping Pan)

Ping reports about implementation experiences with RSVP

Reliable Message
-> soft-state is not useful here
-> reliability problem
-> staged refresh timer (exponential backoff)
-> at the end - very tcp like

Message Packing
-> one session per session
-> size does not matter - but the number of messages do (bundling) (for the cpu)
(extensions allowed juniper routers to support 4x the number of states)

RSVP:

-> cannot handle fragments
-> policy objects go to the limits
-> not an issue for rsvp-te

RSVP does not allow you to transmit the message back from the middle to the initiator.

Two MTU issues:
a) do you get delivery at all (routers/firewalls)
b) loss fragmentation problem (one fragmentation lost)

ad b) efficiency issue

what does a router in the middle of the path do if a RSVP message is fragmented and the fragments do not have the router alert option set?

MPLS
-> for the traffic engineered network there are rarely "route" changes. hence only refresh messages

RFC 2961 is only for refresh messages
does not help with bursts of messages

advantages of using a tcp transport layer:
- only one tcp connection between two nodes
- bundling is better
- mtu discovery
- fragmentation
- ...

MPLS label distribution is no application for NSIS.
RSVP-TE and RSVP is not compatible.

Support for multiple transport layers
- in certain environments you would like to use raw IP, SCTP, TCP or UDP.

Implication of security:

- Using existing security mechanism (Channel security
- Packet size - policy type of objects/pk-based stuff

==========

NSIS interim meeting (Tuesday February 11)

Attendants:

Scott Bradner, John Loughney
Marcus Brunner, Ruediger Geib, Hans Lippitsch, Hannes
Tschofenig, Henning Schulzrinne, Robert Hancock, Kwok Ho Chan, Ping Pan,
Joachim Hillbrandt, Janne Rinne

Agenda:

The Design Space for NSIS signaling protocols (Henning)

My NSIS model

NSIS assumptions:
- we want to support more than a single application
- operating definition of nsis: path associated state management
- possible applications:
  * traceroute
  * active networking
  * network property management/middlebox communication

- functionality which should be supported:
* bi-directional signaling support
* signaling must be done possible between only some nodes (for repair, etc.) / supported in RSVP but only to some extends - no generic mechanism

Finding NSIS peers

we do not want to find all nsis peers along the path
goal : find next node along the path

scott: rsvp makes the assumption that every node along the path supports rsvp

robert: another category is not only to find an nsis node but also to find a node which supports a specific functionality
henning: do we want to provide short-cuts or not?

marcus: explit route object is a externally configured mechanism to avoid discovery (external path knowledge)
scott: the addresses could be any cast addresses

When to discover peers

Who triggers a discovery procedure? (NI, NF)
a) Discovery triggered by NI (requested)
b) Discovery independently at nodes along the network (more traffic but faster reaction time)
e.g. when a route change in the middle is detected.

scott: soft state is important
agreement by everyone - transport connection does not violate this principle

henning:
refresh reduction extension allows you to extend the refresh interval by separating the overloaded functionality of the path message
should the soft state time be configurable?

scott: cbr type of exclusive reservation disallows other nodes to use data traffic
henning: also regular reservation disallows other nodes to make a reservation

henning: automatic discovery is not always what we want. (example: remove state) / this, however, requires further thoughs.

triggered updates in rsvp are only partial useful

three types of route changes
a) observable and meaningful (next nf node is different)
b) observable and not meaningful (possible in an edge model) - not useful to track it
c) non-observable (not immediately)

Next-node discovery

this is the most significant difference between path coupled and path decoupled
the only useful notion for path-decoupled procedure is about the notion of ASes

discovery is an approximation
problems might be caused by load balancing, policy routing, layer 4 load balancing, etc.

there are problems for any type of policy based forwarding and implicit/explicit discovery procedures (path-coupled)

scott: load balancing is even worse since it makes a hash over a fixed portion of the header (including protocol type)

conclusion: let's use the destination address inside the ntlp - load balancing problems always causes problems (i.e. use a simple solution)

Transport Requirements

possible large data volume and large number of messages

henning: minimal nsis application requirements - path -associated state management
scott: "traceroute++" would be an application which is at least required by some people
henning + john: we should not close too many doors to prevent some solution.
john: we should make it very generic
scott: as a summary: there are things that are larger than mtu
henning: see sip as an example

henning: there are two issues - large objects / high bursts
scott: congestion control is very important

Upper/Lower Layers

do all nodes process the messages?
depends what information is located where

scott: extensibility should be provided (see something which i do not know - i should not fail)

scott: see ipv6. some bits which tell what you do when the message is received but not understood

there has been some confusion what the term NTLP means.
henning's definition: ntlp is on top of udp, tcp or some other reliable transport mechanism

Reliability

fast and reliable setup because human factor is involved.

(some issues are only relevant for the initial signaling messages)

Options:
a) end-to-end
* roundtrip estimate is fairly poor
* node processing adds unpredictable variability (e.g. aaa processing).
  (for end-to-end processing you don't have only transmission delay)
b) peer-to-peer
* better rtt
*re-use transport optimizations (sack)
* mandates explicit discovery

Other transport issues

the signaling message can get bigger along the path (add / modify / delete)
designing a new transport protocol is very difficult

Options for transport protocols:
- raw ip
- tcp
- sctp
(see ieee network hol/ sctp)
john: multiplexing capability is the big advantage of sctp not the hol
scott: multiplexing was the reason to create sctp (not hol)
john: sctp has the same order of performance characteristics as tcp
scott: the best things learned from tcp have been added to sctp
further protocols:
- ddp (udp+congestion control)
- dccp

transport protocol issues have no relationship to soft-state principle
optimization issue: state maintenance / state establishment
(local optimization issue whether to keep a transport layer session alive or to close it)

multicast:
even in rsvp nearly all messages are unicast (except for path)
end-to-end principle: not clear what "end" in this context means
henning: nsis is not just message forwarding protocol (like beep) / it is about adding/modifying/deleting objects along the path
scott: it is not modifying the data stream
georgos: is ntlp and nslp always co-located?

State overhead

you are keeping the same information with and without transport layer protocol
transport protocol implementation: mostly a non-issue
transport header overhead: more with transport protocol overhead

Identifiers should be....
(a summary of henning's thoughs)

- we need something
- do not overload identifiers
- should be globally unique (hard todo; an approximation could be sufficient)
- not dependent on host addresses
- should not depend on mac
- constant length
- cryptographically random
scott: an identifier which can be determined / looked up is a bad thing
scott: global identifier requirements are hard to accomplish

henning: there might be more than one identifier

Packet Format

Henning: XML/ASN.1 out of scope (no candidate)
Variations of TLV:

External described:
RSVP: Type has two components (meaning and data type)

Internal described:
Diameter for example

Working Group Update (John)

An application on middlebox communication should be added to the working group
(will be added to the charter)

Requirements: good comments have been submitted / incorporate them and resubmit them to the mailing list

Henning: the requirements do not reflex midcom specific issues

John: My understanding of the requirements draft is help the working group to develop a common language. do not require every possible use. get the requirements draft done - not useful to discuss them too long. henning is requested to review them.

scott: should be real requirements and not desire (less is more)
john: example aaa - many of the required features might not be deployed in future
if a requirement is not in the draft then it does not mean that a solution is not fullfilling

nsis threats should go to last call soon
framework is fairly stable

henning: framework - requirements document inter-relationship (cross-referencing)
robert: framework document

3 critical issues:
- layer split
- mobility issues
- off-path issues

3-6 month a possible target to make final comments

analysis document:

- what have we learned from the past.
- humm at the next ietf meeting to see if the analysis document is useful.
- john thinks there is a need to capture knowledge of what works and what not.
- every working group has to re-investigate security issues (it would be good to have some documents to for example describe current cipher-recommondations)

henning: there are design tradeoff where you never get consensus on. hence it would be better to publish them as a technical report

NSIS Framework Issues (Robert)

Clarification about what the term ntlp means.
inter-layer api?
try to include functions which immediately interact with lower layers

NTLP issues

is there agreement that the ntlp only carries something between two nsis peers?

ntlp has
- to provide discovery
- to interact with routing

upstream/downstream direction:
- ntlp to decide where go
- nslp to trigger it (semantics)

henning: option
- ntlp as a forwarder
- bypass functionality

henning: it is hard to detect a route change in a non-nsis node
a route change is detected downstream -> the upstream can be detected downstream and a message can be transmitted upstream. to provide this functionality nslp functionality can be used.

is reverse routing required by all nslps?
- ni does not care that the operation succeeded/ no message back
- message from the other end could be transmitted back to the ni.

henning: a number problems with that! (security, transport)

assumption: the infrastructure for reverse routing has to be established
if only discovers downstream peers, but must be capable of forwarding message upstream

two basic options for ntlp functionality: Single-hop vs. multi-hop

Single-hop: this would allow an nsis message only be transmitted between two nsis peers

Multi-hop: should it be possible to transmit an empty nslp message over multiple nsis nodes? if so then the ntlp must have more functionality

relationship between middlebox and qos state (for a single session identifier): keep them separate: no strongly

coupling between qos and firewall/nat state - but it should be further discussed on the mailing list.

john: learning from sip is a good thing. a tcp based sip stack is much nicer to work with/implement
henning: state machine got complex due to additional support of udp and because
* semantics of message reception
* message reception
two state machines required

this is only about performance and not about semantics of various messages

scott: multi-stream / hol

scott: congestion control (hop-by-hop or end-to-end)
van jacobson: link-per-link congestion control is not a good thing
henning: congestion control: has to be hop-by-hop
end-to-end is not as possible in the nsis case
henning: flow control at each node and loss less buffering

flow control is the most important property
keeping the edge under network overload conditions in a save state

if we want to have these options (or ever use an existing transport protocol) then we cannot go with end-to-end addressing
john: we cannot decide now which are mandatory/optional

flow identifiers

which information should be included at the ntlp?
the flow identifiers can be arbitrary complex
(just enough to get it through nats)

should wildcarding be allowed. would nats allow this? there was some discussion at the midcom list.

ruediger: an ntlp protocol to itself use a midcom protocol to request a nat binding to subsequently allow the qos signaling protocol go through

henning: holding type of nat state information. hence some solution might not work everywhere.
robert: how much an we do for nats. what is the simplest way to provide this functionality.

henning: can we do anything? we need at least the 5 tuple. the nat box might not be able to keep the binding to long. it is not only about the address re-writing. it is also about requesting a binding lifetime to adjust the binding lifetime

henning: do you allow more than on?
wildcarding issue:
- efficiency (you could transmit more than one message)
- transactional semantics (all or none of them) acit semantics

can nslps update objects in ntlp?

scott: we cannot put any constraints to it.

state management

should the ntlp provide a managment service?

scott: lower latter would create the state installation and management

henning: there is two kind of state
- not transport management type of state
- signaling state information

henning: it does not to be part of the protocol specification. where to locate state is part of the implementation

henning: some state identifier + timer + nslp state which gets logically removed
timer value can be a configured value (specification) or as part of the signaling application

it would be good to describe advantages / disadvantages

john: connection & soft-state is a fuzzy concept

a tcp connection does not need to be related to an NSIS session (one tcp connection might be associated to more than one nsis session instance)
should the ntlp the upper layer nslp that a transport layer connection broke done (might be - api or performance issue)

agreement: ntlp failure does not delete nslp state

john: keep things separate / let the end systems decide / do not combine two different
example: nat and qos signaling
first signaling nat and later signal qos signaling

cedric: this might be expensive

some favor to keep the state information at the nslp

robert: request to provide input on a state management service

Scoping

scott: it cannot be done. there is an endless discussion about this issue on the ipv6 mailing list

if people have well-defined chunks then it should be provided at the nslp.
local objects are also a scoping issues

Rerouting / Mobility events

- how to detect it
- how to handle them (who can handle merging/deletion to avoid double booking)
- is this signaling application specific?

Will be discussed tomorrow

Security issues

should security between non-adjacent nodes be provided?

NSLP NSLP NSLP
NTLP NTLP NTLP NTLP

scott: if the NSIS protocol chooses tcp then it must provide some protection
we don't want add new dos attacks
dos attacks come in various flavors
types of dos attacks and the issues with it should be addressed within the nsis threats document

scott: there are some environments where tcp/sctp might not be the right answer (wireless links)

heavy discussion about transport protocols
henning: we need to decide which properties we would like achieve

robert: would the end-to-end nsis protocol break if a protocol is used between two peers which does not provide a certain property?

robert: mandate requirement at the transport layer
henning: what do you achieve if one hop does not provide reliability in the middle in the network?
henning: the danger is then that people add some functionality for end-to-end layer

georgios: even bob braden specified a raw ip and a tcp type of transport

henning: there are two issues (i am concerned about (b))
a) we do not need certain properties
b) we want to implement at the nslp

scott: dealing with the single link as a special case - hence in a certain environment it might be ok to provide restricted functionality

ACTION: from john to robert: send a mail (what is the current thinking what is needed). anyone should reply with a reason.

scott: is also concerned with tcp in a hop-by-hop environment

Using RSVPv1 as NTLP suggestions for modifications on RFC2205 (Georgios)

ntlp functionality:
- any node should be able to asynchronously generate a signaling message
- load sharing capability
- softstate
- local object
- support for non nsis aware routers along the path

some ntlp should be stateless (skipping the intermediate nodes) - optional / performance enhancement (only used within the admin domain)

henning: do you worry about reliability?
georgios: no - not discussed

robert: couldn't you simply say that you support udp?

scott: we shouldn't rule out a proposals

john: provide additional motivations why it is a good way forward
scott: why isn't it a good way forward

NTLP Design Considerations (Robert)

How does an NTLP entity detect that it is the last one before data receiver?
this issue is caused by the following two reasons:
- new application support (midcom)
- one of the endpoints does not support NSIS

Capability discovery: introduced by providing more generality

A couple of re-route detection mechanisms.

robert: should nsis depend on stun to do the nat business?

scott: let us not do that!

do we optimize for low setup latency?

sven: it is more important to react on changes than a low setup latency

re-routing might be important for some applications (e.g. firewall/nat) - whereas for other applications such as QoS it does not matter too much

in RSVP a refresh is not handled with the same priority as the regular messages

failure management: timers

conclusion: there are new problems introduced - these need to be addressed somehow

==========

NSIS interim meeting (Wednesday February 12)

Attendants:

Scott Bradner, John Loughney
Marcus Brunner, Ruediger Geib, Hans Lippitsch, Hannes
Tschofenig, Henning Schulzrinne, Robert Hancock, Kwok Ho Chan, Ping Pan,
Joachim Hillbrandt, Janne Rinne, and others

Agenda:

Mobility Interaction (Hannes)

Short presentation on the generic limitations of mobility due to the selection of a flow identifier (if the flow identifier is a 5-tuple including the source ip address then a full packet classifier/flow identifier update along the entire path is required).

The session identifier provides some advantages to mobility. Interaction (api, triggers) from the mobility signaling protocol provides performance improvements and is comparable to route detection mechanisms.

Henning: from a protocol design perspective we should support something but an implementation should have the choice to delete something. there are some subtle race condition issues which could cause problems to the entire path

John: a teardown message for the old path might cause some problems - it is a performance improvement
Scott: no symmetric paths can be assumed (if there is more than a single isp along the path then it is very likely that there is no symmetric path anymore).

henning: mobility is likely to introduce significant difficulties.
john: mobility is left for further optimization (not target for the first work)

[marcus] what is the expected time range for a detected path change?
[hannes] depends on: refresh interval, route detection mechanism, interaction with mobility protocol, internal nsis protocol mechanisms (e.g. initiator triggered path detection mechanisms) etc.
[gk] 802.11 (100 feet), doesn't take long to drive 100 feet
[jl] is it 'need to remove state'?
[geib] have to link both reservations
[hannes] one of the observations of rsvp mobileip interaction is the importance of separation of flow identification and session identification
[hs] not only efficiency but mid-session failure for no reason (if you fail
to release resources)
- downlink reservation
* protocol interaction is similar to a route change
[gk] main point: nsis should not discover cross-over router or MAP
[hannes] the discovery of the cross-over router is automatically provided due to existing of the session identifier
[reh] other work suggest you can do it by having a good implementation of
mobility protocol and nsis protocol. could benefit from analysis what type
of messages you send when. but don't think it covers all the cases.
[geib] do all potential anchor points have to be nsis-aware?
[hannes] no.
[marcus] but completely different timing issue: route change is longer
timeframe
[hs] don't care how frequently it happens but rather rapid transition from
old one to new one. cleanup is probably more important for mobility.
[jl] implementation issue?
[hs] issue of old/new path: danger to tear down everything, because
determining the cross-over point is difficult because of race conditions.
[jl] prefer to have it informative. some may want to let the state time out.
[hs] thought agreement this was an optimization. need to provide it but not
mandate the use of it.
[jl] may be some situation where resource reservation is not as 'reserved'
(more advisory than contractual)
[kwok] interdependence between upstream/downstream
[reh] seem to be different, may be dangerous to let one depend on the other.
[kwok] but what about correlation
[sob] will be difficult in the middle of the network
[hs] having coupling between both directions in the network is left for
future work

Wireless and mobility support issues (Georgios)

Generic handover performance requirements

georgios: handover requests could be treated by higher priority
henning: by what? you only need to detect what is new and what is already existing - a policy issue.

minimize number of bits transmitted

Implications using TCP over 2.5 and 3G wireless networks

characteristics of 2..5 and 3g wireless networks

john: this topic is outside the scope of nsis
henning: it effects all transport protocols
john: slow start is a problem
scott: an rfc which addresses this problem will pop-up soon

henning: wireless links affect the performance of protocols using retransmission (it is not a tcp specific problem)
if one can do better than tcp then we should present it.

john: large number of small tcp session might cause problems. we should try to make sure that we avoid
robert: reusing a transport layer connection for a number of message is an option
henning: re-implementing a tcp mechanism within ntlp does not solve the problem.

markus: what of the functions do we need?
john: what are the transport things we need? georgios should provide some of the issues

ietf pilc have studied the impacts and suggest:
- sacks
- tcp timestampe
- roc

Anticipated Handovers (Joachim Hillebrand)

reserving resources in advance
relationship to seamoby
scott: requires movement prediction

john: what in nsis can be done to enable this?
robert: there is a framework issue on this - dependencies are not know

john: joachim should investigate to determine what are the requirements for nsis

henning: a generic mechanism would be good
john: it is currently premature to be done
scott: it is premature to add it to the charter.
john: our work should be mobility be friendly
scott: but it does not require anticipation of the future
scott: there are surely a number of tricky things if there are more than one possible move

NSIS and Mobility - Layer Split & Framework Issues (Robert)

Some issues are open from the framework issues

Is mobility different to re-routing?

reservation theft:
- now system assumes some sort of "return-routability"

home address packet classifier is not feasible
reservation theft is easier

you have to have end-to-end signaling to update the packet classifier

repair is local - there is no option to localize a

possible saving: for the unchanged path you should avoid aaa/policy control

you have to execute the signaling message exchange again
partial nslp execution?

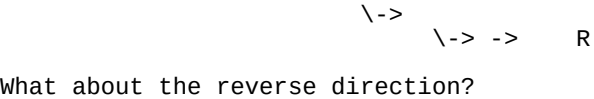

What about the reverse direction?

What about a change in the middle of the middle network?

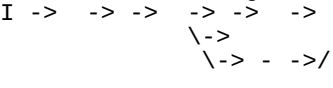

What about a tear-down?

There are some dependencies with aaa/charging and mobility?

Should we support this?
G: for a small change we should
Sven: delegation might be an option

unchanged path requires a packet classifier update

Should we support se-like reservations?
selective wildcard - comes with the possibility of multicast

sven:
- se can only be used between the same points / used in conference type of applications
-

combining different packet classifiers -> merge classifiers and qos
henning: merging on prefixes (this already needs to be done for routing anyway); are routers more flexibility?

detecting crossover routers

- previous / next peer has changed
- are there other issues?

henning:
- session identifier is necessary
- this identifier is sufficient to associate the existing state with an incoming message
- move forward

henning: properties:
- constant over session
- globally unique

ACTION: session identifier was agreed

was is used for?
- detect crossover router
- session ownership problem?

henning: options
- session identifier no security relevant (identity of the owner independent from the id)
- session identifier is security relevant
- session identity (pk-based id)

- globally unique, random collision probably is small, not protected
- nobody can snoop is id (only participating nodes)
- source authentication (authentication, pbk)

markus: decouple security and identifier
henning: done combine too many things with the same mechanism
john: decouple ownership issue + write it down
robert: we have a clear need for it.

henning: move forward - no one came up with a
john: state in the document - it cannot be used as a ownership proposal
john: overloading the id is a bad idea as henning used

ACTION: currently cryptographically random id (not protected - no ownership)

CT/CARD issues

ACTION: john - keep it open
action item for john to look at 5.3.5 of the framework and to comment on it.

cedric: session id for group of flows and different nslp
henning: no; separate issues (a) flows (b) nslp

ad (b) we haven't made this decision yet.

what is the semantically tying the two nslp instances (foo/bar) - what means the combination of foo and bar -> foo fails but bar is allowed etc. this gets very complicated and we don' t want to do it.

john: this sounds like an optimization which should not be considered yet; this is not an ntlp discussion

ACTION: ad (a) it should not be done

henning: there are many things would like to add / keep the core spec small enough - otherwise it will never get implemented.

we don't have to get them into the *-00 version.

john: we should not try to do 110 % - instead we do 90 % and then we have look at deployment/implementations. extensions can be added later.

robert: the capability should be simple

cedric: Capability negotiation required?

henning: you have to provide something for the discovery. there are different type of issues: there are some things that have mandatory, discoverable and optional parts. the optional parts can cause an error if not understood or blindly forwarded. every good protocol has this type of functionality.

henning: there is a danger: if there are two modes and you have a full fletched tp and you have a tp which provides nothing for specific links. bad: functionality implemented in two places (interactions get complicated)

if you do what rsvp refresh reduction did then every transport person would laugh us out of the room. it is not a solution - it is a hack.

ruediger: there has to be something mandated - otherwise the interoperability is in danger!

john: what are the options?

NSIS Meeting Summary(John)

henning: milestones?

john: have been updated
henning: next due date?
john: requirements to IESG
nsis threats quite ready
meta-question for analysis doc

john: mobility stuff should be captured. possible rolled into the analysis document.

henning: lessions are always in the eye of the beholder.
john: we don't want to spend more time on the analysis document

sven: off-path signaling bof next time
(1 hour session)

john: describe why it is needed for a restricted environment.
henning:
a) people have a different notion of what off-path is? how you scope the problem.
b) can be quite complicated

john: restricted to intra-domain signaling (inter-domain can always get very complicated)

sven: it might not only be a qos application
john: one or two good use-cases might be a good thing