[manet] Reactive routing protocols, what are the differences?

22 messages

Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> Thu, Nov 1, 2012 at 11:22 AM
To: "manet@ietf.org" <manet@ietf.org>
Cc: "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

The obviously best people to answer this should be document authors, but anyone else may have useful additions and comments. Ideally the different document authors could agree a list. (If they differ in that one has X and the other doesn't, but one wants to say "we plan to add/remove X" then X should be listed as a difference with that caveat, in at least my ideal world.)

If we set aside, for the moment (though these things matter):
- The presentational quality of the documents,
- Any issues of 5444 compliance and other formatting issues,
- Issues of internal data organisation,
- Minor details such as possible different timeout parameters etc.
then what are the technical (and I stress that word) differences between DYMO and LOADng? (I'm withholding the term AODVv2 for reasons I may come back to.)

Note that it's a lot more useful to have direct differences than differences of each from AODV (especially when both have the same difference). And it would be useful to have the objective differences separated from the "and now why this is better" discussion - though that would be a next step.

I'm not saying I don't see any of the differences. But I certainly haven't worked out the complete list. In trying to form my view of how things should go forward (a view that is coming together, and when it does, I'll argue for it) and I hope for other people as well, it would be good to know what the differences are. Regardless of views for or against each, we should be able to objectively list the significant differences - if we can't then something is wrong.

--
Christopher Dearlove
Senior Principal Engineer, Communications Group
Communications, Networks and Image Analysis Capability
BAE Systems Advanced Technology Centre
West Hanningfield Road, Great Baddow, Chelmsford, CM2 8HN, UK
Tel: +44 1245 242194 | Fax: +44 1245 242124
chris.dearlove@baesystems.com | http://www.baesystems.com

BAE Systems (Operations) Limited
Registered Office: Warwick House, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hants, GU14 6YU, UK
Registered in England & Wales No: 1996687

****
This email and any attachments are confidential to the intended
recipient and may also be privileged. If you are not the intended
recipient please delete it from your system and notify the sender.
You should not copy it or use it for any purpose nor disclose or
distribute its contents to any other person.
****

_____

manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Abdussalam Baryun <abdussalambaryun@gmail.com> Thu, Nov 1, 2012 at 12:34 PM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Hi Chris,

I think that your suggestions and other discussions related to the subject MUST be done on the MANET list, and not within the f2f meetings.

AB
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Ulrich Herberg <ulrich@herberg.name> Thu, Nov 1, 2012 at 12:54 PM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Chris,

I think this is a good start to have a technical discussion. I hope that everyone could thoroughly *read* both documents and give a technical opinion. I will read the latest DYMO revision in detail and then answer to your email.

Best regards
Ulrich

On Thu, Nov 1, 2012 at 9:22 AM, Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> wrote:
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Abdussalam Baryun <abdussalambaryun@gmail.com> Thu, Nov 1, 2012 at 1:37 PM
To: Ulrich Herberg <ulrich@herberg.name>
Cc: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>, "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Thanks Ulrich, I agree and would like to know your opinion on this subject and other co-authors so we can have a progress discussion on the manet-list. I will do the same.

AB

[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Ulrich Herberg <ulrich@herberg.name> Thu, Nov 1, 2012 at 5:02 PM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Hi Chris,

you have seen my review on DYMO. I will try to answer to your question, and focus on the technical differences, not presentation.

On Thu, Nov 1, 2012 at 9:22 AM, Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> wrote:
> The obviously best people to answer this should be document authors, but anyone else may have useful additions and comments. Ideally the different document authors could agree a list. (If they differ in that one has X and the other doesn't, but one wants to say "we plan to add/remove X" then X should be listed as a difference with that caveat, in at least my ideal world.)
>
> If we set aside, for the moment (though these things matter):
> - The presentational quality of the documents,
> - Any issues of 5444 compliance and other formatting issues,
> - Issues of internal data organisation,
> - Minor details such as possible different timeout parameters etc.
> then what are the technical (and I stress that word) differences between DYMO and LOADng? (I'm withholding the term AODVv2 for reasons I may come back to.)

First, let's see what is common. Both are reactive protocols, using RREQ, RREP and RERR. So if someone claims that DYMO performs great and LOADng badly in the same scenario, I cannot understand that. MANET has understood the scenarios where reactive protocols are useful and where not.

Now, to the differences:

- DYMO cannot be end-to-end secured. Messages are changed in transit (and not just hop-limit or the metric), but rather addresses can be removed from RERRs and RREQs. There is also no provision to allow external mechanisms to add additional reasons to reject messages as invalid.
- DYMO uses the originator address in an address block, LOADng in the message header. The sequence number is a TLV value in DYMO, and LOADng uses the message sequence number. DYMO requires the originator address to be the first one in the address block, the destination must be the second one. LOADng uses a TLV to determine the target address.
- DYMO can advertise multiple addresses in an RERR; they can be removed in transit of the message.
- DYMO allows intermediate routers to reply (as an option). That makes end-to-end security difficult. In the core DYMO, there is a destination sequence number that may be contained in RREQs in DYMO.
- DYMO allows for unicast RREQ, but does not specify in detail how to use that.
- There are four timers for each route entry in DYMO, only one in LOADng.
- LOADng can be used on other layers; DYMO is tied to IP.
- LOADng provides a bidirectionality verification using RREP_ACK, a

time-out of these, a blacklisted set and a Pending Acknowledgment Set to verify bidirectional links. DYMO says that other mechanisms can be used, but does not specify these.
- DYMO has several options for expanding ring RREQ, precursor list, adding route information in transit, message aggregation in RFC5444 packets and reporting multiple unreachable addresses in a RERR. LOADng takes the approach to have a slim core of a basic mechanism that is applicable in all MANET use cases, and companion documents with extensions. In DYMO, it is not clearly specified what happens if some routers support an option, and others don't.
- LOADng uses a Metric message TLV, and it is clearly defined how to update the metric under way. If a router in transit does not recognize a route metric type, it is reset to a "hop count" tlv extension type of the Metric TLV and the value set to 0xFFFFF... (for the full TLV length). It is specified that security mechanism must ignore the content of the metric TLV value and that the length cannot be changed under way, so that end-to-end security is possible. DYMO uses an optional "distance" field for the metric, which is not clearly specified how it is updated. Also, since this is optional, it is unclear if routers receiving a message and forwarding it, update the distance field or not.
- LOADng allows for (optionally) waiting to reply with a RREP, in case a "better" RREQ comes a little later. In DYMO, a RREP is always sent immediately.

There are probably more differences, but I let other chime in.

Best regards
Ulrich
[Quoted text hidden]

Mukul Goyal <mukul@uwm.edu> Thu, Nov 1, 2012 at 10:46 PM
To: Ulrich Herberg <ulrich@herberg.name>
Cc: "Christopher Dearlove (UK)" <Chris.Dearlove@baesystems.com>, manet@ietf.org, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Hi Ulrich

Thanks for pointing out the key differences between the two protocols. Going over this list and having read both drafts, I have the following opinion:

1) It seems to me that there are no technical reasons why the two drafts cannot be merged. The fact that it was not possible to merge the two proposals is unfortunate.
2) Perceived strengths of LOADng, such as facilitating end-to-end security, can easily be adopted in DYMO/AODV2. Bidirectionality verification can be done either inside the protocol or in an independent manner.
3) It seems to me that, functionality wise, LOADng is a restricted version of DYMO/AODVv2. In other words, it is possible to configure a DYMO deployment to behave like a LOADng deployment. Charlie made the same point in an earlier message. If LOADng is chosen in place of DYMO/AODVv2, we will lose nice features you pointed out:

"- DYMO has several options for expanding ring RREQ, precursor list,
adding route information in transit, message aggregation in RFC5444
packets and reporting multiple unreachable addresses in a RERR."

These features seem useful in general MANET scenarios.

Just my $0.02.

Thanks
Mukul
[Quoted text hidden]

JP Vasseur (jvasseur) <jvasseur@cisco.com> Fri, Nov 2, 2012 at 3:12 AM
To: Mukul Goyal <mukul@uwm.edu>
Cc: "Christopher Dearlove (UK)" <Chris.Dearlove@baesystems.com>, "<manet@ietf.org>" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Hi Mukul,

On Nov 1, 2012, at 11:46 PM, Mukul Goyal wrote:

> Hi Ulrich
>
> Thanks for pointing out the key differences between the two protocols. Going over this list and having read both drafts, I have the following opinion:
>
> 1) It seems to me that there are no technical reasons why the two drafts cannot be merged. The fact that it was not possible to merge the two proposals is unfortunate.
> 2) Perceived strengths of LOADng, such as facilitating end-to-end security, can easily be adopted in DYMO/AODV2. Bidirectionality verification can be done either inside the protocol or in an independent manner.
> 3) It seems to me that, functionality wise, LOADng is a restricted version of DYMO/AODVv2. In other words, it is possible to configure a DYMO deployment to behave like a LOADng deployment. Charlie made the same point in an earlier message. If LOADng is chosen in place of DYMO/AODVv2, we will lose nice features you pointed out:

JP> This is exactly the point that I made "several" emails ago. Charlie proposed to include in DYMO/AODVv2 functionalities of LOAD-ng,
with options. This is the option 1.
Fully agreeing with your analysis.

Thanks.

JP.
[Quoted text hidden]

Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> Fri, Nov 2, 2012 at 5:15 AM
To: Abdussalam Baryun <abdussalambaryun@gmail.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

You may notice that I asked this question on the list. But there is no either/or. Things are done on the list. Then when we have a meeting things are done there. And after the meeting things are done on the list. For the formal process, read the appropriate RFCs.

--

Christopher Dearlove

Senior Principal Engineer, Communications Group
Communications, Networks and Image Analysis Capability
BAE Systems Advanced Technology Centre
West Hanningfield Road, Great Baddow, Chelmsford, CM2 8HN, UK
Tel: +44 1245 242194 | Fax: +44 1245 242124

chris.dearlove@baesystems.com | http://www.baesystems.com

BAE Systems (Operations) Limited
Registered Office: Warwick House, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hants, GU14 6YU, UK
Registered in England & Wales No: 1996687

From: Abdussalam Baryun [mailto:abdussalambaryun@gmail.com]
Sent: 01 November 2012 17:35
To: Dearlove, Christopher (UK)
Cc: manet@ietf.org; Thomas Heide Clausen (thomas@thomasclausen.org)
Subject: Re: [manet] Reactive routing protocols, what are the differences?

***** WARNING *****

This message originates from outside our organisation, either from an external partner or the internet.
Keep this in mind if you answer this message.
Please see this process on how to deal with suspicious emails.

[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> Fri, Nov 2, 2012 at 5:22 AM
To: Ulrich Herberg <ulrich@herberg.name>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

There is, superficially at least, an apparent contradiction between your points:

- DYMO cannot be end-to-end secured. Messages are changed in transit
(and not just hop-limit or the metric), but rather addresses can be
removed from RERRs and RREQs.

which implicitly suggests LOADng does not do this

and

- LOADng uses a Metric message TLV, and it is clearly defined how to
update the metric under way.

Could you expand on this please?

--
Christopher Dearlove
Senior Principal Engineer, Communications Group
Communications, Networks and Image Analysis Capability
BAE Systems Advanced Technology Centre
West Hanningfield Road, Great Baddow, Chelmsford, CM2 8HN, UK
Tel: +44 1245 242194 | Fax: +44 1245 242124

chris.dearlove@baesystems.com | http://www.baesystems.com

BAE Systems (Operations) Limited
Registered Office: Warwick House, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hants, GU14 6YU, UK
Registered in England & Wales No: 1996687

-----Original Message-----
From: Ulrich Herberg [mailto:ulrich@herberg.name]
Sent: 01 November 2012 22:02
To: Dearlove, Christopher (UK)
Cc: manet@ietf.org; Thomas Heide Clausen (thomas@thomasclausen.org)
Subject: Re: [manet] Reactive routing protocols, what are the differences?

----------! WARNING ! ----------
This message originates from outside our organisation,
either from an external partner or from the internet.
Keep this in mind if you answer this message.
Follow the 'Report Suspicious Emails' link on IT matters
for instructions on reporting suspicious email messages.

[Quoted text hidden]

Jiazi YI <ietf@jiaziyi.com> Fri, Nov 2, 2012 at 5:53 AM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Hi Chris,

I think Ulrich is in deep sleep at the moment, so please allow me to have some words on this.

For reactive protocols, updating the route information (hop-count, metric) is inevitable. In the specification of DYMO, the messages can be changed relatively arbitrarily, including removing addresses from the messages. The intermediate RREP also makes end-to-end security impossible.

LOADng clearly defines which fields in the routing messages can't be changed, and which fields are mutable. For example, for RREQ:

> The following fields of an RREQ message are immutable, i.e., they MUST NOT be changed during processing or forwarding of the message: RREQ.addr-length, RREQ.seq-num, RREQ.originator, and RREQ.destination.
>
> The following fields of an RREQ message are mutable, i.e., they will be changed by intermediate routers during processing or forwarding, as specified in **Section 12.2** and **Section 12.3**: RREQ.metric-type, RREQ.route-metric, and RREQ.hop-count.
>
> Any additional field that is added to the message by an extension to this protocol, e.g., by way of TLVs, MUST be considered immutable, unless the extension specifically defines the field as mutable.

This allows the protocol to secure the messages by zeroing the mutable fields.

best

Jiazi

(sorry to Chris if you received multiple copy of this message. I fixed one typo though :) My previous one was bounced by manet mailing list because of not using the right sender address)

[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> Fri, Nov 2, 2012 at 6:42 AM
To: Jiazi YI <ietf@jiaziyi.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Having anything other than hop limit and hop count mutable is not the security mechanism suggested in 6622/5444.

I'm of the view that the current approach of securing NHDP, securing OLSRv2 etc. is not where things should ideally be, that the ideal place is at the 5444 multiplexer wherever possible. If doing hop by hop security, then that is the place, as it's where packets live. But if we could do it once for all message types, then that's a major gain.

Now as soon as LOADng has anything else mutable, that doesn't help that. OK, it's better than nothing, but those fields are buried in TLVs (using 5444) and messy.

I'm at a disadvantage, I haven't studied why LOADng needs those fields (other than hop count) mutable - or if it really does. But it reduces "here's a clear advantage over DYMO" to "here's a more partial advantage over DYMO". And I think Ulrich's summary could be edited to better present this.

So if we adopted my separate proposal, this would be on the menu: why are those fields mutable? Do they have to be? Can we find an alternative? (Unfortunately, I can see why probably not. But it's still a question.)

[Actually I can see an alternative, which is a much more limited metric, which takes small integer values, and we increase hop count not by one but by this metric, limiting paths to maximum 255 metric. We could still get hop count from hop limit if we knew how it started - e.g. in a non-mutable TLV. But I strongly doubt this is

good enough.]

--

Christopher Dearlove

Senior Principal Engineer, Communications Group
Communications, Networks and Image Analysis Capability
BAE Systems Advanced Technology Centre
West Hanningfield Road, Great Baddow, Chelmsford, CM2 8HN, UK
Tel: +44 1245 242194 | Fax: +44 1245 242124

chris.dearlove@baesystems.com | http://www.baesystems.com

BAE Systems (Operations) Limited
Registered Office: Warwick House, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hants, GU14 6YU, UK
Registered in England & Wales No: 1996687

From: Jiazi YI [mailto:yi.jiazi@gmail.com] **On Behalf Of** Jiazi YI
Sent: 02 November 2012 10:54
To: Dearlove, Christopher (UK)
Cc: Ulrich Herberg; manet@ietf.org; Thomas Heide Clausen (thomas@thomasclausen.org)
Subject: Re: [manet] Reactive routing protocols, what are the differences?

***** WARNING *****

This message originates from outside our organisation, either from an external partner or the internet.
Keep this in mind if you answer this message.
Please see this process on how to deal with suspicious emails.

Hi Chris,

[Quoted text hidden]
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Henning Rogge <hrogge@googlemail.com> Fri, Nov 2, 2012 at 7:01 AM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

On Fri, Nov 2, 2012 at 12:42 PM, Dearlove, Christopher (UK)
<Chris.Dearlove@baesystems.com> wrote:
> Having anything other than hop limit and hop count mutable is not the
> security mechanism suggested in 6622/5444.

Thats my interpretation of RFC6622 too. I want to implement RFC6622 (at least packet and message level ICVs) in my generic "PacketBB" reader/writer core. As soon as more fields become mutable this would get very messy at best.

> So if we adopted my separate proposal, this would be on the menu: why are
> those fields mutable? Do they have to be? Can we find an alternative?
> (Unfortunately, I can see why probably not. But it's still a question.)

I am not convinced that securing a Distance Vector Protocol "end to end" is as easy as doing it with a Link State Protocol.

Its the nature of Distance Vector Protocols that they Nodes exchange their private idea how to whole world works with their neighbors (local exchange of global information). Link state protocols only exchange their local data, but they flood them to the network (global exchange of local information), which makes things easier security wise.

Securing everything except the metric will make it impossible to pretend to be a node that isn't in the mesh, but you still could spoof your local distance to the node in any way you want.

Maybe using Address Block ICV TLVs would help, they allow the protocol to explicitly state what should be included into the signature. This could even include parts of the message header.

(Does BGP have some ideas how to do end-to-end security? Its the only Distance Vector Protocol I know that is widely deployed.)

Henning Rogge
--
Steven Hawkings about cosmic inflation: "An increase of billions of billions of percent in a tiny fraction of a second. Of course, that was before the present government."
[Quoted text hidden]

Teco Boot <teco@inf-net.nl> Fri, Nov 2, 2012 at 7:10 AM
To: Henning Rogge <hrogge@googlemail.com>
Cc: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>, "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Op 2 nov. 2012, om 13:01 heeft Henning Rogge het volgende geschreven:

> (Does BGP have some ideas how to do end-to-end security? Its the only
> Distance Vector Protocol I know that is widely deployed.)

We have RIP*, Cisco has EIGRP.
In many deployments, routers have shared secret. I don't see much gain in end-to-end security, as control plain info passes middle routers.

Teco
[Quoted text hidden]

Jiazi YI <ietf@jiaziyi.com> Fri, Nov 2, 2012 at 7:12 AM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Hi,

The reason that LOADng needs those fields mutable is to support different metrics other than hop-count, even routers using different metrics in the same routing domain can at least find a path.

To support different metric types, a metric-type tlv and a route-metric tlv are needed. The route-metric tlv has to be mutable because it needs to be updated at each hop.
Having metric-type tlv mutable can make supporting different metrics in the same network possible.

It's common that in a heterogenous network, there are various transmission medias (802.11, 802.15.4, cable, PLC...), therefore possible different metrics.
In LOADng, if a router A (with metric-A) gets an RREQ message that it doesn't understand (say, metric-B), router A will change the metric-type to HOP_COUNT, and forward the message (the hop-count field is always used). For the destination of RREQ, it will first consider the metrics that it understands, and then HOP_COUNT. Of course, this will result in "degrading" to HOP_COUNT for certain routes, but at least we can get a usable route.

I'm just introducing the design of LOADng, and would appreciate any good idea on this issue.

best

Jiazi
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> Fri, Nov 2, 2012 at 7:17 AM
To: Teco Boot <teco@inf-net.nl>, Henning Rogge <hrogge@googlemail.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

In a link state protocol there is a real gain. (I'm afraid I've not got enough time right now to discuss why.) For a reactive protocol which relies not just on the end message, but on the intermediate sequence that the RREPs pass through, that's not so obvious. At least not without an accumulating signature (which the maths exists for without signature size growth, I believe). Those could be a 6622/5444 option. It's an interesting area of discussion if we move into the technical issues.

--
Christopher Dearlove
Senior Principal Engineer, Communications Group
Communications, Networks and Image Analysis Capability
BAE Systems Advanced Technology Centre

West Hanningfield Road, Great Baddow, Chelmsford, CM2 8HN, UK
Tel: +44 1245 242194 | Fax: +44 1245 242124
chris.dearlove@baesystems.com | http://www.baesystems.com

BAE Systems (Operations) Limited
Registered Office: Warwick House, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hants, GU14 6YU, UK
Registered in England & Wales No: 1996687

-----Original Message-----
From: manet-bounces@ietf.org [mailto:manet-bounces@ietf.org] On Behalf Of Teco Boot
Sent: 02 November 2012 12:10
To: Henning Rogge
Cc: Dearlove, Christopher (UK); manet@ietf.org; Thomas Heide Clausen (thomas@thomasclausen.org)
Subject: Re: [manet] Reactive routing protocols, what are the differences?

----------! WARNING ! ----------
This message originates from outside our organisation,
either from an external partner or from the internet.
Keep this in mind if you answer this message.
Follow the 'Report Suspicious Emails' link on IT matters
for instructions on reporting suspicious email messages.

[Quoted text hidden]

****
This email and any attachments are confidential to the intended
recipient and may also be privileged. If you are not the intended
recipient please delete it from your system and notify the sender.
You should not copy it or use it for any purpose nor disclose or
distribute its contents to any other person.
****

[Quoted text hidden]

Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> Fri, Nov 2, 2012 at 7:32 AM
To: Jiazi YI <ietf@jiaziyi.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

I understand the basics. And I agree I can't see (except my bracketed aside) how to do it without a mutable field. But having that mutable field reduces the value of the argument against DYMO from "DYMO is mutable, LOADng is not" to "DYMO is (uncontrolled?) mutable, LOADng is managed mutable". I'm not convinced by the argument about having to mutate metric type. If you can't rely on it being available to your network layer protocol consistently in the one MANET, heterogeneous though it may be, I'm not sure you have a well-put-together network. You could always make the metrioc increment when unknown the maximum such value.

But there is, as another post raised, the larger question of what end to end message authentication buys you. If in a route A-B-X-C-D, where X is a bad guy, if X relays all RREQs and RREPs flawlessly, but throws all data packets on the floor, X has done his job, and without needing to forge anything. You also need B and/or C to authenticate X. Lower layer? (in which case why can't that do the whole job?). RREP-ACK? Accumulating signature? Something else?

(Note that a link state protocol gets the B-X and X-C done. Those are, after all, links.)

--

Christopher Dearlove

Senior Principal Engineer, Communications Group
Communications, Networks and Image Analysis Capability
BAE Systems Advanced Technology Centre
West Hanningfield Road, Great Baddow, Chelmsford, CM2 8HN, UK
Tel: +44 1245 242194 | Fax: +44 1245 242124

chris.dearlove@baesystems.com | http://www.baesystems.com

BAE Systems (Operations) Limited
Registered Office: Warwick House, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hants, GU14 6YU, UK
Registered in England & Wales No: 1996687

From: Jiazi YI [mailto:yi.jiazi@gmail.com] **On Behalf Of** Jiazi YI
Sent: 02 November 2012 12:13

[Quoted text hidden]

[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Teco Boot <teco@inf-net.nl> Fri, Nov 2, 2012 at 8:00 AM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Op 2 nov. 2012, om 13:32 heeft Dearlove, Christopher (UK) het volgende geschreven:

> I understand the basics. And I agree I can't see (except my bracketed aside) how to do it without a mutable field. But having that mutable field reduces the value of the argument against DYMO from "DYMO is mutable, LOADng is not" to "DYMO is (uncontrolled?) mutable, LOADng is managed mutable". I'm not convinced by the argument about having to mutate metric type. If you can't rely on it being available to your network layer protocol consistently in the one MANET, heterogeneous though it may be, I'm not sure you have a well-put-together network. You could always make the metrioc increment when unknown the maximum such value.
>
> But there is, as another post raised, the larger question of what end to end message

> authentication buys you. If in a route A-B-X-C-D, where X is a bad guy, if X relays all RREQs and RREPs flawlessly, but throws all data packets on the floor, X has done his job, and without needing to forge anything. You also need B and/or C to authenticate X. Lower layer? (in which case why can't that do the whole job?).

There could be 2nd order nodes: hosts. Or the routing protocol security mechanism is implemented in the routing deamon, which is very similar to the first.

> RREP-ACK? Accumulating signature? Something else?
>
> (Note that a link state protocol gets the B-X and X-C done. Those are, after all, links.)

There can be malicious routers with link state protocols too.
Also, there is a requirement that all routers share the same policy on what to do with failed authentication, and result of checking must be the same on all nodes. Not easy to deploy. And missing in olsrv2 core protocol.

Teco
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Dearlove, Christopher (UK) <Chris.Dearlove@baesystems.com> Fri, Nov 2, 2012 at 9:01 AM
To: Teco Boot <teco@inf-net.nl>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

> Yes, of course there can be malicious routers with link state protocols. I've even helped implement one. But everything used is a link, and all links are carried in messages/packets, and if all messages/packets are authenticated, then all links are authenticated, and hence no malicious routers are included.
>
> But the RREQ/RREP process (at its simplest) uses information that is not included in any message.
>
> Going back to OLSRv2, yes, all routers will have to agree on a process. But that's another discussion that I won't get into right now.

--

Christopher Dearlove

Senior Principal Engineer, Communications Group
Communications, Networks and Image Analysis Capability
BAE Systems Advanced Technology Centre
West Hanningfield Road, Great Baddow, Chelmsford, CM2 8HN, UK
Tel: +44 1245 242194 | Fax: +44 1245 242124

chris.dearlove@baesystems.com | http://www.baesystems.com

BAE Systems (Operations) Limited
Registered Office: Warwick House, PO Box 87, Farnborough Aerospace Centre, Farnborough, Hants, GU14 6YU, UK
Registered in England & Wales No: 1996687

From: Teco Boot [mailto:teco@inf-net.nl]
Sent: 02 November 2012 13:00
To: Dearlove, Christopher (UK)
Cc: Jiazi YI; manet@ietf.org; Thomas Heide Clausen (thomas@thomasclausen.org)

[Quoted text hidden]

[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Ulrich Herberg <ulrich@herberg.name> Fri, Nov 2, 2012 at 10:00 AM
To: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>
Cc: "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen (thomas@thomasclausen.org)" <thomas@thomasclausen.org>

I am glad that we finally have some technical discussions...

I agree that it is more difficult to provide end-to-end security in a reactive protocol than in a proactive link-state. The idea in LOADng is that we know exactly the fields that are mutable, and we know that the length will not change nor the position of any of the fields (but potentially the value of the metrics tlv). So it is relatively easy to just zero out the Metric TLV value field. But I admit it's less straight forward than in OLSRv2.

Ulrich
[Quoted text hidden]
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Christopher Dearlove <christopher.dearlove@googlemail.com> Fri, Nov 2, 2012 at 10:51 AM
To: Ulrich Herberg <ulrich@herberg.name>
Cc: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>, "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen(thomas@thomasclausen.org)" <thomas@thomasclausen.org>

But that still leaves the problem that with just end to end security, the RREP path is not reported in any message and hence is not protected, and my A - B - X - C - D example.

--
Christopher Dearlove

christopher.dearlove@gmail.com (iPhone)
chris@mnemosyne.demon.co.uk (home)
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Jiazi YI <ietf@jiaziyi.com> Fri, Nov 2, 2012 at 11:14 AM
To: Christopher Dearlove <christopher.dearlove@googlemail.com>
Cc: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>, "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen(thomas@thomasclausen.org)" <thomas@thomasclausen.org>

I agree that in your A-B-X-C-D example, end-to-end security is not enough.
But it can protect the network from certain attacks. For example, it can prevents X flooding numerous RREQs, or X spoofing the identify of D to send a RREP.

And I think your comments on metric-type TLV make sense. I'll discuss with other LOADng authors.

best

Jiazi
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet

Ulrich Herberg <ulrich@herberg.name> Fri, Nov 2, 2012 at 11:28 AM
To: Christopher Dearlove <christopher.dearlove@googlemail.com>
Cc: "Dearlove, Christopher (UK)" <Chris.Dearlove@baesystems.com>, "manet@ietf.org" <manet@ietf.org>, "Thomas Heide Clausen(thomas@thomasclausen.org)" <thomas@thomasclausen.org>

Chris,

I agree with you that this is an issue the WG has to work on. For now, I don't have answer to that, but I would like to discuss that with you and other interested people in Atlanta and on the list.

Best
Ulrich
[Quoted text hidden]

_____
manet mailing list
manet@ietf.org
https://www.ietf.org/mailman/listinfo/manet