3. Ballot Comments

NAME Jim Burns
COMMENT TYPE E
CLAUSE 5.2
PAGE 34
LINE 1
COMMENT START
This sentence indicates that "An implementation of a MAC Security Entity (SecY) for which full conformance to this standard is claimed shall not implement Cipher Suites other than those specified in Clause 14.", but then under section 5.2 there is a statement "c) Use Cipher Suites not specified in Clause 14, but meeting the criteria specified in 14.2, 14.3". These two statement seem contradictory. Shouldn't the optional capabilities be in addition to the required?
COMMENT END
SUGGESTED CHANGES START
Change the statement in clause 5.1 to be "An implementation of a MAC Security Entity (SecY) for which full conformance to this standard is claimed shall not implement Cipher Suites other than those specified in Clause 14 or allowed by the criteria in 14.2, 14.3."

SUGGESTED CHANGES END

this is for “conformance”

NAME John Viega
COMMENT TYPE TR
CLAUSE 5.4
PAGE 34
LINE 18
COMMENT START
I believe we should be strict about ciphers and modes that can be approved. I'd like to get somewhat technical in the text here. My suggestions may be considered too draconian.
COMMENT END
SUGGESTED CHANGES START
The use of additional cipher suites must meet the following guidelines

1) The underlying cryptographic ciphers must be endorsed either by NIST or the NESSIE standards project.
use a different word than endorsement, approval

2) The cipher suite must provide message authentication using a message authentication algorithm with a academically peer-reviewed proof of security against forgery attacks, even in a model where the attacker has the ability to choose messages for the sender.

3) If confidentiality is provided, the confidentiality mechanism must have an academically peer-reviewed proof of security in a model where the attacker has the ability to adaptively choose both plaintexts and ciphertexts.

4) Mechanisms for confidentiality and message authentication must be used in a way that is consistent with their proof of security. For instance, if using the CBC mode of operation, the IV must be randomly selected with each message, and not sequential.

5) If serviced by separate algorithms, the properties of the authentication and confidentiality mechanisms must be combinable in accordance with well-established security results. Either the encryption must happen before authentication, or the encryption must be performed through keystream generation.

6) move to beginning. strength Algorithms chosen must have an effective key length of at least 128

bits.
In schemes built on block ciphers, the underlying block cipher must have a block width of at least 128 bits. no known attacks with complexity <2 ^100 work

see Mick's comments on structure of conformance clause

NAME Jim Burns
COMMENT TYPE T
CLAUSE 6.7
PAGE 40
LINE 47
COMMENT START
The text in this section indicates '...if stations are added to the CA, MAC_Operational transitions to False in either all the stations originally participating in the CA or in all those added, ...]'. For implementation it will be necessary to specify how to choose the group that shall transition MAC_Operational to False. Otherwise there shall be interoperabiltiy issues that will result in both group transitioning MAC_Operational False (wasting time) or neither group transitioning MAC_Operational False (causing a security issue). As the CA is 'invisible' to the SecY, this shall presumably occur at the discretion of the KaY.
This clarification should also occur in section 7.2, p51, line 28.
COMMENT END
SUGGESTED CHANGES START
Add a sentence "Determining which group shall transition MAC_Operational to False is outside the scope of this specification and shall be defined within IEEE 802.1af and signaled through the LMI."
SUGGESTED CHANGES END
specification.. and is defined in IEEE802.af, see Mick's comments on this clause as well. too much information here.

NAME Paul Congdon
COMMENT TYPE TR
CLAUSE 6.7
PAGE 41
LINE 1
COMMENT START
adminPointtoPointMAC does not take on the value of 'TRUE' as is implied here. It is either ForceTrue, ForceFalse or Auto. The algorithm for all the choices needs to be better specified. The MAC_Operational you are talking about is the lower ISS MAC_Operational as well, not the one that MACSec is trying to drive.
COMMENT END
SUGGESTED CHANGES START
I think you want adminPointtoPointMAC to be set to auto and at most one in the CA for this to work. I supposed it could be set to forceTrue as well, but this should require changes to the current definition of adminPointToPointMAC. If it is set to forceFalse, then operPointToPointMAC must be false regardless of the number of stations in the CA.

SUGGESTED CHANGES END:
it's only auto that is important
if ForceTrue always true, ForceFalse, always false

NAME Tony Jeffree
COMMENT TYPE ER
CLAUSE 6.9
PAGE 42
LINE 21-24
COMMENT START

The Editor's Note clearly needs to be removed; however, it highlights the fact that right now we don't have any formal means of recording maintenance items for 802.1D.

COMMENT END
SUGGESTED CHANGES START

Remove the Editor's Note.

Need to discuss what to do with the note otherwise - i.e., how we plan to record/action ongoing maintenance of 802.1D.

SUGGESTED CHANGES END
back to Mick and Tony

NAME Paul Congdon
COMMENT TYPE T
CLAUSE 6.10
PAGE 43
LINE 7
COMMENT START
Actually, MACSec across a provider bridge network runs the risk of increasing the amount of frame loss due to replay protection in the presence of frame re-ordering. It might be possible to see frames re-ordered across a provider network due to prioritization or internal link aggregations. If replay protection is on, the amount of frame loss could go up dramatically
COMMENT END
SUGGESTED CHANGES START
This is a good place to document the issue of replay protection enabled across a provider network and how that could increase frame loss. Also, in general, while 802.1 tries to minimize frame re-order, there is a chance and if replay protection is enabled, frames that would have normally been delivered out of order would now be dropped. Insert some sentences with the essence of the above text.
SUGGESTED CHANGES END:
Accept. purpose of section to highlight dilemmas. Paul will input suggestions.

NAME John Viega
COMMENT TYPE T
CLAUSE 7.1
PAGE 46
LINE 11
COMMENT START
Realistically, SCs are going to be limited to 2^64 octets or so, given the current scheme (and use of AES). I think this is what is meant when the note mentions "many years without interruption", but it might be good to add an explicit number.

Also, it's worth noting that, as long as the scheme uses a single root symmetric key, this is probably the practical limit, before you need a new key that is randomly chosen and distributed in some out of band method.
COMMENT END
SUGGESTED CHANGES START
Not sure... this may be worth discussing.
SUGGESTED CHANGES END

many years refers to fact can will use succession of new master keys
lifetime of sngle symmetric key is 2^64. when derivation of new key is from an old key, what's the "information leakage". but completely fresh master keys obviates this concern. lifetime of a series of keys related to a single master key would be problematic.
this should be said somewhere, probably not in this section.

getting an entirely fresh symmetric key.

NAME: Allyn Romanow
COMMENT TYPE: TR
CLAUSE: 7.1
PAGE: 48
LINE: 27
COMMENT START:
It's not absolutely clear from the draft whether a port is allowed to accept non-SecTAGged packets while it is in a CA. Clearly, it does support control
frames from other EtherTypes, for example 802.1X, but it's unlikely that it will accept data frames from outside the CA.
The text says “While D can send and receive frames using the insecure connectivity provided by the shared LAN, it does not have SAKs that would allow it to participate in any of the SAs that currently support SCA, SCB, or SCC..”

This sounds like D can communicate with the members of the CA, and probably it cannot. In any case this needs clarification.

Another relevant section is, Section 8.2, which says
The KaY will set the NeighborsAllSecYs variable if every adjacent station has a SecY.
COMMENT END:
SUGGESTED CHANGES START:
If D cannot communicate with the members of the CA, the text should say something like
Members of the CA will not accept packets from non-members.

Also, the document should be checked for other references to communication between members and non-members
of the CA

SUGGESTED CHANGES END:

don’t confuse stations and ports. where validateframes is strict, D can’t send to member of CA and will not appear in ControlledPort.

Not in this section, station D can send to Uncontrolled ports in A and B

NAME John Viega
COMMENT TYPE TR
CLAUSE 7.1.2
PAGE 49
LINE 22
COMMENT START
I think this text and the graphics previous to this are going to confuse people. My current understanding is that all SCs share a single symmetric key, and the SC is more about nonce selection. If this is the case, I think we should say that there is generally only one SC that all participants use for transmitting and receiving, because otherwise this will continue to be an ongoing source of confusion.
COMMENT END
SUGGESTED CHANGES START
Depends on the resolution, but I'll be happy to provide text.

john- text doesn’t allow one shared SC.
Mick- doesn’t want to preclude SC per transmitter

should be made clear that the keys do not have to differ

NAME Paul Congdon
COMMENT TYPE ER

CLAUSE 7.1.3
PAGE 50
LINE 17
COMMENT START
How does the SecY know it has all the keys it needs? I think the case being talked about here is one where the MAC_Operational was once TRUE and everything was fine, then all of a sudden there were no keys because they aged out and weren't replaced by the KaY in time. It would be worth mentioning how a SecY can get into this state.
COMMENT END
SUGGESTED CHANGES START
Include a statement that this case can occur after the CA is up and running. Include the conditions that could cause the SecY to not have the keys it needed.
SUGGESTED CHANGES END:

Historically, originally the SecY had more knowledge of it's own state.
it's the KaY that drives this.
keep idea of this, talk about the KaY
the Kay will drive mac_oper_false when run out of PN and KaY has gone to sleep.
SecY knows when PN is exhausted.
best SecY can do, on xmit if out of PN, mac oper comes down
on recv, goes false if can't recv from any, how know? no recv SA in use,

NAME: Allyn Romanow
COMMENT TYPE: TR
CLAUSE: 7.1.3, 9.6
PAGE: 50, 66
LINE: 10, 21
COMMENT START:
The text is not consistent as to the number of SAs that must be stored by a receiving station.
p.50 line 10 says receiver has to store 3 SAs
p.66, line 21, cl 9.6 says a receiver needs to support 2 SAs
COMMENT END:
SUGGESTED CHANGES START:
Change p. 50 to
capable of storing SAKs for [three] two SAs for each inbound SC,
And check the doc for any other inconsistent references to the number of required SAs per SC.
SUGGESTED CHANGES END:

3 were in case new master at precisely the same time as change SAK, would cause extra time to get Master Key. Onn objected. should allow extra time in unlikely case. so went to 2 keys.

NAME Paul Congdon
COMMENT TYPE TR
CLAUSE 7.2
PAGE 50
LINE 35
COMMENT START
I believe multiple instances of a CA are possible on a single LAN by using the SCI to demultiplex and/or look-up the instances. There does not need to be multiple common ports to achieve this. A single common port will do, but the look-up function on SCIs needs to change to allow this. The text argues that some other form of multiplexing is required (e.g. EPON LLID, etc), but it is possible to do this using the SCI.
COMMENT END
SUGGESTED CHANGES START
Reword much of this clause pending the discussion and presentation of the multiple-CA material at 930 on 3/15/05.
SUGGESTED CHANGES END:
NAME Mick Seaman
COMMENT TYPE T
CLAUSE 7.2

PAGE 50
LINE
COMMENT START

The discussion of multiple service instances in this clause is now a lot technically weaker than it was. I know this was and still is an issue for a number of people but that does not mean that the same idea should be repeated in the document as may times as possible, nor is repetition of observations required. Saying the same thing in multiple different ways simply means there are more sources of inaccuracy to correct. The extent of the changes to this clause are not justified by the disposition of comments on D2.0 (I have checked).

In the first paragraph (pg 50, line 37), it is not true that it is not possible to have multiple Common Ports from a single ISS, it would just be that they would get you exactly the same thing - so would not necessarily produce multiple instances. It is further not true that there can be only one SecY attached to a single LAN - since there can be different SecYs in different systems. The paragraph is making the mistake of trying to conduct a tutorial at exactly the same time as the basic facts are being laid down, so consequences of multiple decisions are misrepresented as consequences of a single fact, or as straight forward assertions. Further the term "Common Port" is not introduced until clause 10, so use of it in definitive text causes a dependency that cannot be properly satisfied in a document that has to have a linear order. Similarly use of the term SecY to mean anything particularly definite should be avoided in Clause 7.

I summarize the suggested changes below (after SUGGESTED CHANGES), but I think it is worth describing how they are assembled, step by step. Given the confusion and dispute that can be caused by inaccuracy I have tried for as much accuracy as possible. In particular I have made the distinction (glossed over in the rest of the text, and let us keep it that way, because it just leads to text expansion and nothing more) between a service instance, which is properly a connectionless association (supported by necessary protocol, including its identification) and a access point for (or point of attachment to) that service instance. A (service) access point is how an entity attaches to a service instance. The names "Controlled Port", "Uncontrolled Port", and "Common Port" are labels for service access points. Thus it can be seen that the sentence fragment "it is not possible to have multiple Common Ports from a single ISS" could have been precisely interpreted as "it is not possible to have multiple Common Ports for a single service access point for an instance of the ISS" which is more precisely stated as "it is not possible to have <multiple service access points <for an instance of the ISS>> for a <single service access point for an instance of the ISS>" (angle brackets inserted to parse the sentence) which reduces to "an object A is not the same thing as multiple instances (greater than one) of object A", i.e. as saying nothing new at all.

The first sentence of the first paragraph should remain, it can be improved by the insertion of "service access point for an instance of the" (which is sufficiently precise to get over the problem described immediately above) with similar supporting changes. The first part of the second sentence was imprecise and described above, and is now no longer required. The second part is also wrong in detail as previously described, so the second sentence should go entirely.

The second paragraph is actually more restrictive than absolutely logically necessary (or can be read as such with the lack of precision involved in using "instance" instead of "access point for instance"), which will get us into trouble with some ways of supporting multi-access LANs. Moreover there could be multiple Common Ports without multiple instances of the insecure MAC service. When I tried to make the existing text more precise I found that the first and second sentence ended up saying exactly the same thing, with a change in word order. Using the slightly more compact text in D2.0

(which was the base of the second paragraph) avoids this problem and leads to

"Multiple instances of the secure MAC Service can be provided by a single LAN provided that each instance
is uniquely identified by unencrypted fields contained in each received frame. These fields identify separate
instances of the unsecured MAC Internal Sublayer Service, each capable of supporting a distinct service access point for each of a number of SecYs."

These two sentences can be added to the end of the first paragraph, where they logically belong.

The third paragraph is unnecessarily restrictive, just being true most of the time, and should be deleted. It also repeats information that is in the fourth paragraph (after the long NOTE), the first sentence of which in turn duplicates information in the second paragraph. I don't think a networking savvy audience needs to be explicitly told that fields in a frame that allow sets of frames to be distinguished compose a multiplexing function, and if this information is put immediately after the first paragraph with nothing in between it doesn't have to repeat information in that paragraph. The allusion to Provider Bridges also needs to be made more specific. This allows the fourth paragraph to be simplified.

The long NOTE 1 was originally part of a comment that I submitted on D2, but not part of the suggested replacement text. It is far too long and casual for standard text. Clearly the ideas need capturing in the document, but what is required is a definite recommendation (should) rather than a NOTE. This text should appear after, and not before, the ideas currently in the fourth paragraph (as changed above).

COMMENT END
SUGGESTED CHANGES START

Replace the first four paragraphs (i.e. those before NOTE 2) and NOTE 1 of 7.2 with the following

"
Each service access point for an instance of the secure MAC Service is supported by a service access point for an instance of an insecure MAC Internal Sublayer Service. Multiple instances of the secure MAC Service can be provided by a single LAN, provided that each instance is uniquely identified by unencrypted fields contained in each received frame. These fields identify separate instances of the unsecured MAC Internal Sublayer Service, each capable of supporting a distinct service access point for MAC Security.

Identification of each insecure service instance, and multiplexing and demultiplexing to and from the transmission capabilities provided by the LAN, can performed wholly below the ISS by a media specific or media dependent functions. Some media are defined to support such a multiplexing function, e.g. the LLID used by P802.3ah EPON (See Clause 12). Provider Bridges are also capable of supporting multiple instances of the ISS over a network of individual LANs (See 11.6).

MAC Security should not be used to support multiple instances of the secure MAC Service on a single physical LAN without the use of unencrypted frame fields to identify separate instances of insecure service, each supporting a single instance of secure service. While the use of security to provide multiplexing is impossible to prevent (since different cryptographic keys can be used to separate connectivity) relying solely on security to define the connectivity makes deployment and fault management difficult - the topology of an entire network could change as security was enabled or disabled on a single LAN. Key agreement protocols that use the insecure MAC

service can require a matching instance of that service for each secure service instance.

NOTE 1-The service access point for the secure MAC Service is referred to as Controlled Port of the MAC Security Entity (SecY, Clause 10) and the service access point for the insecure MAC Service as the SecY's Common Port. Access to the insecure service for protocol entities above MAC Security is provided at the Uncontrolled Port.
"
SUGGESTED CHANGES END
multi-access
tbd how to treat in .1AE, incorporate or have a separate doc
go over last 2 paragraphs- Mick and Paul

NAME Dan Romascanu
COMMENT TYPE TR
CLAUSE 8.1.7
PAGE 58
LINE 20
COMMENT START
It is not clear what is the design requirement related to Intrusion Detection. The first phrase in the text seems to say that the management function can facilitate intrusion detection, while the second phrase makes a claim about detecting abnormal traffic patterns which is not substantiated by any details (like what counters?)
COMMENT END
SUGGESTED CHANGES START
Delete this section.
SUGGESTED CHANGES END

NAME Mick Seaman
COMMENT TYPE E
CLAUSE 8.1.7
PAGE 58
LINE 22-31
COMMENT START

If anything is to be said here it needs to be more definite, and the reference provided. The use of "might" indicates suspect text which will be a target in later ballots.

COMMENT END
SUGGESTED CHANGES START

Replace the text of this clause with

"Intrusion detection is facilitated by integrity and replay protection, and the management counters (10.7) that record the receipt of invalid (presumably modified) and repeated and misordered (likely to be replayed) frames. Management for client policies (7.3) that use the guaranteed connectivity provided by MACsec should also record attempted violations."

Delete the two editor's notes.

SUGGESTED CHANGES END
delete
counters signify abnormal behavior

NAME Dan Romascanu
COMMENT TYPE TR
CLAUSE 8.2.4
PAGE 60
LINE 14 and following
COMMENT START
The requirement in this clause seems to contradict the non-goal q) in Section 1.2, which defines

discovery of relationship between peers as a non-goal of the standard
COMMENT END
SUGGESTED CHANGES START
delete this section, or non-goal q) in Section 1.2
SUGGESTED CHANGES END

NAME Mick Seaman
COMMENT TYPE E
CLAUSE 8.2.4
PAGE 60
LINE 14-34
COMMENT START

This clause and those following are written as they were a normative clause for the KaY, which can't be because the document is about MACsec not the KaY. It also incorrectly uses the word "must". It is somewhat out of date as the topics touched upon are now covered in clause 10, and contains a number of small hints/notes to the author as to KaY design which can now be taken out. Clearing up the appearance of being normative etc. should be handled by making definitive statements ("is" rather than "must", "shall", "will" etc.)
.

COMMENT END
SUGGESTED CHANGES START

In the first para, replace "must be able to discover" with "discovers". Delete the second sentence.

In the second para replace "must accept" with "accepts". Delete the following two sentences.

In the third para replace "must accept" with "accepts", and "will deliver" with "delivers". Delete the last (bracketed)sentence.

Delete the fourth para (single sentence).

SUGGESTED CHANGES END

NAME Jim Burns
COMMENT TYPE T
CLAUSE 8.2.6
PAGE 60
LINE 48
COMMENT START
This section indicates "The KaY provides authorization of services to be delivered to a peer station based on the outcome of the authentication and authorization process." The previous section (8.2.5 p 60 line 39) it indicates "In this case, the key management process will find a pre-shared key and operate without the authentication process needing to generate the key". The question is, if there is no authentication process where does the authorization come from? Presumably it is a policy within the station.
COMMENT END
SUGGESTED CHANGES START
Change sentence on line 48 in section 8.2.6, p 60 to only reference authorization
"The KaY provides authorization of services to be delivered to a peer station based on the outcome of the authorization process. This authorization process is based on the policies of the station and the context of the connection which may include authentication."
SUGGESTED CHANGES END
truncate at to a peer station.
rationale- don't want to hamstring .1af

NAME Mick Seaman
COMMENT TYPE T

CLAUSE 8.2.7
PAGE 61
LINE 3-22
COMMENT START

See my comment on 8.2.4. Remove interesting asides that are out of scope as well, such as first para and reference to Master Key. The last sentence of the third para is just flat wrong as it does not conform to the model for SAs (point to multipoint) already explained. Comments on policies and their coupling to authorization are out of this scope - they are controlled by key agreement not the SecY.

COMMENT END
SUGGESTED CHANGES START

Delete the first para (line 3/4).

Replace "will deliver", "will create", "will accept", with "delivers", "creates", "accepts" whenever they occur.

Delete the last sentence of the third para (line 10), and replace "SCs" with "SCs and SAs" in the prior sentence.

Delete the second sentence of the fourth para (lines 13/14).

Delete all but the first sentence of the last para (line 19/20).

SUGGESTED CHANGES END
accepted
==========END MONDAY

NAME Karen Randall
COMMENT TYPE ER
CLAUSE
PAGE
LINE
COMMENT START
I'm uncomfortable with approving this given the state of the document -- there are still sections to be completed. This document seems to be a little premature to be circulated for full working group ballot.
COMMENT END
SUGGESTED CHANGES START
The document needs to be cleaned up and empty sections completed.
SUGGESTED CHANGES END

NAME Karen Randall
COMMENT TYPE ER
CLAUSE 3.22
PAGE 20
LINE 45
COMMENT START
Strengthen the definitions by incorporating definitions from other security standards, where appropriate.
COMMENT END
SUGGESTED CHANGES START
Modify the current definition of nonce to incorporate the definition from the X9F standards (given in X9F TR1)

A non-repeating value, such as a counter, used in key management protocols to thwart replay and other types of attack.

SUGGESTED CHANGES END

NAME Karen Randall
COMMENT TYPE ER
CLAUSE 3.23
PAGE 20
LINE 34-36
COMMENT START
Strengthen the definitions by incorporating definitions from other security standards, where appropriate.
COMMENT END
SUGGESTED CHANGES START
The definition of non-repudiation from the X9F standards (given in X9F TR1) is

This security service provides proof of the integrity and origin of data - both in an unforgeable relationship - which can be verified by any party.
SUGGESTED CHANGES END

NAME Dennis Volpano
COMMENT TYPE ER
CLAUSE 6
PAGE 35
LINE 53
COMMENT START
Authentication and authorization is outside ...
COMMENT END
SUGGESTED CHANGES START
Replace "is" with "are"
SUGGESTED CHANGES END

NAME Frank Chao
COMMENT TYPE E
CLAUSE 6.5
COMMENT START
any default value for adminPoint2PointMac ?
COMMENT END
SUGGESTED CHANGES START
Provide the default values.

NAME Ken Patton
COMMENT TYPE T
CLAUSE 6.10
PAGE 43
LINE 6
COMMENT START

The text does make clear how the MACsec service will of necessity provide
^^^
NOT

a lower effective MTU than the unencrypted MAC layer will provide. Since there is a "tax" of SECtag headers to be paid, then the effective MTU offered by the MACsec service will always be less than then MTU of underlying media, even as the MTU of the media (such as an expected increase in 802.3 frame size) grows arbitrarily huge. Since Annex Z.5.4 states that MACsec will not pursue fragmentation, implementors must be made aware that the header tax will impinge on the frame size of the payload.

COMMENT END

SUGGESTED CHANGES START

Add additional language specifying the expectation that MACsec's effective MTU is lower than the MTU of the unencrypted media.

SUGGESTED CHANGES END

NAME Les Bell
COMMENT TYPE T
CLAUSE 7.3.1
PAGE 52
LINE 40, 52-54
COMMENT START
Bullet (b) and Note 2 describe a VLAN classification that is not supported in other sections of the document. For example, there is nothing said on how to associate a VLAN ID to a CA, SC, or SA.

This also applies to the last paragraph on page 53.
COMMENT END
SUGGESTED CHANGES START
Discuss whether this is intended and, if so, how this VLAN classification is configured and how it inter-operates with the PVID, protocol-based VLAN classification, the 802.1Q VLAN Tag, and the 802.1ad VLAN Translation Table.
I suggest that MACsec is not used for VLAN classification purposes.
SUGGESTED CHANGES END

NAME Les Bell
COMMENT TYPE T
CLAUSE 8.2.4
PAGE 60
LINE 14-15
COMMENT START
The definition of the Discovery mechanism, whether it is a protocol or not, and whether it uses the Bridge Group Address, is a matter for the P802.1af standard.
MACsec should constrain itself to stating the requirements the KaY must meet to be compatible with MACsec.
COMMENT END
SUGGESTED CHANGES START
Replace the last sentence with "The Discovery mechanism must be constrained to peer stations on an individual LAN."
SUGGESTED CHANGES END

NAME Michael Wright
COMMENT TYPE T
CLAUSE 8.2.4
PAGE 60
LINE 15
COMMENT START
The discovery mechanism is in question.
Should P802.1ab be cited or is this outside of the project?
COMMENT END
SUGGESTED CHANGES START
If P802.1ab is the correct mechanism cite it else state the discovery mechanism is out scope.
SUGGESTED CHANGES END

NAME Michael Wright
COMMENT TYPE TR
CLAUSE 8.2.5

PAGE 60
LINE 38 & 43
COMMENT START
Line 38 says the KaY may authenticate Line 43 states that SecY assumes that authentication has occurred. This seems inconsistent to me.

COMMENT END
SUGGESTED CHANGES START
If the SecY assumes authentication then the KaY should always due authentication else SecY should not assume that authentication has occurred.

SUGGESTED CHANGES END
accept. subclause needs clarification

NAME Mick Seaman
COMMENT TYPE E
CLAUSE 8.3, Figure 8-2
PAGE 61
LINE 53/54
COMMENT START

The concept of MACsec AAD was introduced in attempt to clearly specify the boundary between decision within MACsec and choices left up to specification of the cipher suite support of MACsec. The idea was to keep as much of the application specific detail away from the cipher specification part as possible. Unfortunately this approach has not worked well, and with reasonable options for some cipher suites to protect the PN and SCI as part of their IV, rather than as "AAD" it looks as if the idea of "MACsec AAD" has just served to complicate rather than simplify. Attempts to clarify have resulted in it becoming less rather than more precise, so it needs to be removed.

COMMENT END
SUGGESTED CHANGES START

Delete the last sentence on page 61. Remove "MACsec AAD" from Figure 8-2.

SUGGESTED CHANGES END
proposed- accept

NAME Les Bell
COMMENT TYPE TR
CLAUSE 8.3
PAGE 62
LINE 13
COMMENT START
The validation function takes the Secure Data as an input and returns the User
Data.
COMMENT END
SUGGESTED CHANGES START
Replace "the octets of the Secure Data are returned" with "the octets of the User Data are returned".
SUGGESTED CHANGES END

NAME: Allyn Romanow
COMMENT TYPE: ER
CLAUSE: 9.2
PAGE: 64
LINE: 2
COMMENT START:
The text isn't clear about whether the ICV field should be 16 octets or 8 to 16 octets.

The field length of ICV in Figure 9-1 says 8 to 16.
in 9.11, the text says
“The length of the ICV is Cipher Suite dependent, but is not less than 8 octets and not more than 16.”
However, other places in the text refer to the ICV as 16 octets.
It seems preferable to have the field fixed at 16, and if a cipher suite wants to use less, it can pad the rest of the field.
COMMENT END:
SUGGESTED CHANGES START:
Change either the text that suggests the field is variable or the text that says the field is fixed at 16 octets
SUGGESTED CHANGES END:

NAME Paul Bottorff
COMMENT TYPE T (Technical)
CLAUSE 9.8
page 66
COMMENT START
At 10GE the re-keying time will be about 40 minutes. This may be too quick.
COMMENT END
SUGGESTED CHANGES START
Reconsider PN field to extend re-keying time.

SUGGESTED CHANGES END

NAME Glenn Parsons
COMMENT TYPE T
CLAUSE 9.8
PAGE 66
LINE
COMMENT START

In section 9.8, a 32-bit packet number field is introduced. At
10gb/s, with maximum length packets, that's roughly 42 minutes between re-key events.

COMMENT END

SUGGESTED CHANGES START

If re-keying at intervals less than every few hours is a problem, then we need to re-think the PN field. IPSEC had to deal with this, and now has an ESN (Extended Sequence Number) scheme to support larger replay spaces, thus reducing the re-key frequency.

SUGGESTED CHANGES END
The re-keying at 10G is every 5 minutes. This poses absolutely no issue for hardware processing. It has to generate a 128-bit random number, do AES encryption and send a couple of packets.

NAME: Allyn Romanow
COMMENT TYPE: ER
CLAUSE:9.5
PAGE:66
LINE: 3-9
COMMENT START:
The use of the C bit is not made sufficiently clear. The name "changed" seems to cause confusion.
COMMENT END:
SUGGESTED CHANGES START:
<Re-write the text.>
Since the mandated ciphers do not change the text, it would be less confusing
to rename this field to something like “Reserved for use by alternate cipher suites”.
SUGGESTED CHANGES END:

NAME Jim Burns
COMMENT TYPE TR
CLAUSE 9.9
PAGE 67
LINE 10
COMMENT START
We use the value 00-00-00-00-00 as a special SCI. It is my understanding from some issues that occurred in 802.11 that the 00-00-00 OUI is owned by Xerox (but not used). Do we require permission from Xerox to use this value?
COMMENT END
SUGGESTED CHANGES START
Determine if we require permission from Xerox to utilize the 00-00-00 OUI.
SUGGESTED CHANGES END

NAME Les Bell
COMMENT TYPE T
CLAUSE 9.9
PAGE 67
LINE 2
COMMENT START
The SCI does not provide replay protection.
COMMENT END
SUGGESTED CHANGES START
Remove bullet (c).
SUGGESTED CHANGES END

NAME Frank Chao
COMMENT TYPE E
CLAUSE Figure 10.5
page 79
COMMENT START
In the upper left hand corner of the flow chart, sa->next_PN = rx.pn + 1 ; update_lowest_pn (next_PN, replayWindow), it may cause the replay window moves backward and forward.

COMMENT END
SUGGESTED CHANGES START
Misk suggested it should be changed to
sa->next_PN = max(rx.pn + 1, sa->next_PN) ; update_lowest_pn (next_PN, replayWindow);

where the max() function returns the greater of its two arguments.

SUGGESTED CHANGES END

NAME Dennis Volpano
COMMENT TYPE T
CLAUSE 10.6.2
PAGE 77
LINE 26
COMMENT START
What is preliminary replay detection?
COMMENT END
SUGGESTED CHANGES START
Include the replay detection that uses a window in Annex Z as part of replay detection described in this clause, and eliminate "preliminary".
SUGGESTED CHANGES END

NAME Les Bell
COMMENT TYPE TR

CLAUSE 10.6.3