IEEE 802.1 Minutes, May 2004

Interim Meeting, Monday AM, May 17, 2004

Opening Remarks – Tony Jeffree

Meeting Arrangements – Dolors Sala

1:30 break for lunch

9:00 am start

11:00 am break

Administrative stuff – Tony

Voting membership

WG and TG operations

Patent Policy – Tony Jeffree

The required two slides were shown and Tony insured that the folks in the room are aware of the IEEE patent policy

September Interim – Tony Jeffree

Plan has fallen apart because ITU-T does not allow meeting fees so can not have a joint meeting

Liaison reports – Tony Jeffree

Need to consider ITU liaison letters this week.

Task Group Schedule for the week – Mick and Dolors

802.1ag PAR issues – Tony Jeffree

Amendment to 802.1Q NesCom said can not open a new PAR because already have a revision to 802.1Q so 802.1ag will revise the revision of 802.1Q

If more than two PARs after two years/three then need a new PAR to revise the base standard

Need to decide what will be in the new base revision

There will be some discussions this week to determine what will be in the new base. Things that may/could be in revision

Sorting priority scheme stuff

This needs to happen after 802.1ad

Need a plan about how to handle this

TIA TR41.4 Work on LLDP Extensions – Paul Congdon

Media Endpoint Detection

http://www.ieee802.org/1/files/public/docs2004.....

Enable deployment of specific security policies

Refining scope

Requesting liaison with 802.1

With liaison, 802.1 may review and vote on the document

Corporate membership

Need to get a list of voting members and corporate reps so folks in 802.1 can know who is representing which company

Disposition of ballot comments 802.1ab – Paul Congdon

Zero no votes, 44 comments

Mick's comment 14 End of LLDPDU TLV has to be mandatory on transmission

Bit of discussion about Mick’s comment 19 to make sure conservative on transmit but liberal on receive

Discussion about send and receive rules for the end of LLDPDU TLV

Comment 18 – Karl Weber

TTL TLV with length greater than 2

Comment 42 – Dan Romascanu

Comment 43 – Matt Squire

Only send no error so this is not a useful TLV

Remove this from the 802.3 TLV

Paul has talked to the folks with significant comments and all have agreed to save for sponsor ballot. The editorial comments will be incorporated and a note made that certain comments (Mick’s and Matt’s) will be made against the sponsor ballot draft.

Currently, we do not have conditional approval from the exec to run a sponsor ballot.

Also, we need to allocate missing values in the document before running the sponsor ballot

There will be a re-circulation ballot and then a sponsor ballot after the July plenary

Summary of LinkSec work for the week

IEEE LinkSec, 802.1AE and 802.1af, and IEEE 802.1X, May 17-20, Barcelona

Most of the discussion was on IEEE 802.1af, Key Management, with little time spent on IEEE 802.1AE MAC Security.

IEEE 802.1AE - A considerably updated version, D2.0, released just prior to the meeting, with a Task Group ballot for June 15. Therefore, not much discussion- that will be for next meeting in July. Note that the model of Security Associations for MACsec is not the same as that for IPsec, see Clause 7 for MACsec's model.

Some emphasis on how MACsec and KEYsec interact, who does what. Also, emphasis on migration and deployment. Attention to turning on security in a step by step fashion, making sure each step works.

IEEE 802.1X - the ballot has passed, however there were some comments to address. If all goes well, draft will go to Rev Con in September and if approved, it will become a standard.

IEEE 802.1af - There were several presentations describing people's thoughts on issues. Still very preliminary in nature. Decision that .1af will not do authentication, but will assume authentication has already taken place through some other mechanism such as pre-shared keys, SNMP, or public key exchange. IEEE 802.1af will provide protocol for key management and authorization. A presentation described the requirements for the LMI Layer Management Interface between the MACsec and KEYsec elements.

Overview of 802.1ae D2.1 – Allyn Romanow

- There is a task group ballot open until June 15
- Disposition of comments for D1.2 is on the web site
- All comments have been resolved except for Dan's jitter comment
- Re-organization of Material – most of the material is the same but the document has been re-organized
- Keys- Key nomenclature – master key used for entire session from which a short term key is derived
  - Need to store 3 Secure Association Keys (SAK)
  - Number of messages to derive new SAK should be 0
  - Discussion of the relationship of Key Agreement and the Key Hierarchy
- E bit doesn't say whether encryption or not, says whether there's been anything appended to the field, or you can parse the field
- Interoperability and Migration
  - Got rid of Null Cipher Suite and Include Tag – reduces unnecessary complexity
  - Now use management controls to control whether you are doing confidentially – E bit is bit 3 of TCI

Further overview of 802.1ae D2.1 – Mick Seaman

- EPON – Single copy broadcast SCB
  - It is possible to spoof OLT since this is symmetrical key encryption, that is, there may be confidentially but not integrity
  - There are ways to allow broadcast confidentially and create a separate key for integrity but this is value add. This could be a lot of complication at this level
  - Using SCB because it's better to use 1 bit to say it's an SCB
- Management
  - MAC operational parameters are described in terms of CA and SA
  - There is a problem with this draft, which is MAC operational – MAC can receive and may be able to transmit. The receivers have to be alive before any transmitters are started, otherwise there is a startup problem – this still needs fixing.
- SecY Management parameters
  - Need to reflect how real systems can get stuck
  - There may be a hole between encoding headers and encrypting such that counters bounce the wrong way
  - After this ballot we should be to the point of creating state machines so this can be resolved
- MACsec Operation
  - Couple of new diagrams
  - Should show that the input on the transmit side matches the input of the receive side such that things are encrypted and decrypted correctly
  - KaY is like another user of the uncontrolled port
- SecY Operation

MACsec migration – Mick Seaman

On the website http://www.ieee802.org/1/files/public/docs2004/SecureRstp02.pdf

Norm - using MACsec to protect control protocols before you had h/w to protect data. The doc flushes this out. You can stop unauthorized entities from changing the root; etc.

Deployment challenges

Current draft is not complete but it is starting to put the structure together to allow for step wise deployment

What is in a deployment plan – should enable you to step back easily, small steps can rollback easily, learn something constructive from each step. From each deployment step want positive feedback before going forward. Try to determine what might go wrong

What services is your network providing? Requirement - want to not have anything peculiar going on when put in security

Want to selectively turn on security.

Standard management controls necessary.

Migration step by step

Step 1 to step 2 - everyone that should be managed is managed.

This can be determined by the untagged count not incrementing

Step 3 only receive tag frames, so any entity that does not have MACsec deployed will not be able to communicate

Transmit tagged frames; transmit and receive tagged frames

Step 4 makes sure the key agreement protocols are up and running

Step 5 check replay validation count. At this point things should be up

Monday PM, May 17, 2004

IEEE 802.1X Ballot disposition – Tony Jeffree

The ballot has passed, however there are comments to address, and 3 "no" voters whose comments need to be addressed.

Mostly editorial comments

Cl 6.7 - bi-directional, mutual authentication, using the clause worked out on mailing list

Comments from Adrian Stephens asking for considerable architectural clarification. It was decided, after discussion, that it would require an extensive amount of work to do this, and would be considered at the next revision of the document.

Cl 7.3 Will add format for 802.11, in place the format for 802.5 which is there now.

Spent time on Jonathan Edney's comments, mostly accepted and changed draft.

There will be a recirculation ballot. If all goes well, the draft will go to Rev Con in September. If they approve it, it becomes a standard. They recommend to standards board.

Compact GVRP – Norm Finn

The presentation is on the web site at http://www.ieee802.org/1/files/public/docs2004/Compact-GVRP-slides.zip

Goals

Reduced number of packets to transmit 4k VLAN states

  - Reduce number of timers
  - Maintain full compatibility with standard GVRP
- One PDU carries complete GVRP state
- Reduce the number of timers
  - If everything fits in one PDU only need one timer
- Point to Point links
  - Use the timer a bit differently to optimize the number of control messages that must flow to the remote
  - Still need anti-chatter
  - What is the default timer value for point to point?
    - Order of one second
    - Discussion about the timers
      - There are historical issues with timers considering FDDI and 10 Mbps Ethernet
      - All of these assumptions can be changed but need to consider that many HW gets a one second timer tick
      - Could have different timers for point to point versus shared
      - Could have different timers for link speed
  - How do you know this is a point to point link?
    - Can not be for sure
    - But can assume that only two of you know this protocol on this link
    - GVRP is not always turned on in this case there is management controls to turn on the required VLAN
    - There may be a race condition
      - If more than one party then startup may have problem
        - Nope, if lost the state machine timeout will handle this case
  - Maintain shared media compatibility
  - Maintain compatibility with GVRP
    - Is it really important?
      - Can an old version supply VLAN?
      - Do we care?
        - Probably don't care.

A Multiple VLAN Registration Protocol (MVRP) – Mick Seaman

- http://www.ieee802.org/1/files/public/docs2004/MVRP-Introduction-030.pdf
- How to localize topology changes over a specific tree
- Constraining topology changes to a specific part of the network is a bit more difficult than first glance
- When physical topology changes then topology changes then MVRP asked for VLAN if and only if it really needs it
- There is a risk of false positives but probably not of false negatives
- If there is interest in doing this then speak up to see if this should be put in the standard
- The consensus of the room is this is something to consider
- This is too elaborate to go into the Q par

- This needs some thought about where it goes
- Don’t allow this to hold up .1ad

Simulation Report – Paul Bottorff

- EGVRP Basic Concepts
- Domain with size associated with a given domain with a default of 12 bit

Discussion about scaling and VLAN tag size – Mick Seaman

- 12 bit is where we are
- Another view is this is a new thing so figure out how many bits can be sustained
- Are these multiple set of 12 bits or a single set?
- Can scaling be achieved by more VLANs or by using more boxes?
- How to move forward?
- Changing the number of bits has ramifications that must be considered – byte counters and VLAN counters must increase
- Customers today want per port per VLAN counters so the increase gets difficult
- Can not increase the number of VLANs and the amount of stuff a VLAN has to be aware of
- Rebuttal – 4k VLANs was sufficient for enterprise but it is not sufficient for the provider networks. The current deployment is point to point in the provider space so the 4k limit is a problem
- Some different ideas to work our way past this issue
  - What is the real problem?
  - Like know if any one thinks 4k service instances are enough? Nope, no one is say this
  - When talking about VLAN address space is it really bits in the packet?
  - MPLS is not the only way to wire up a set of bridges
  - Clearly a solution where provider bridges connect up into some mesh
  - When sending a packet have to consult the address function to determine where to send the packet
  - If ten thousand customers then would need ten thousand ways to determine which customers
  - Hierarchical is a way to think about the problem but the questions is 12 bits enough
  - You need lots of service instances and a service instance will consume at least one VLAN
  - What is the restriction of the places the service instance must be mapped to a VLAN?
  - What we are specifying – what is going across the boundary to the customer from the service provider?
  - In one physical location support 50K services
    - At some point you are creating a wire don’t want to look at VLAN until it is a small number
  - 4k VLANs is not enough for service instances
  - Whole lot of ways to expand the number of service instances MAC in MAC, MPLS, etc;
  - We should talk about some of these ways to get by this impasse
  - Agree there must be some type of hierarchy

With the fact several folks are shipping Q in Q then we need a standard for this but there is probably a better way that we need to work through

A technical requirement - The encapsulation must not create a situation where packets flow to places they should not

The independence of the customer traffic must be guaranteed by what ever solution is developed

Tuesday AM, May 18, 2004

Connectivity Fault Management met separately

Thoughts on KeySec – John Viega

- In Orlando, seemed to agree on .af phases:
  - Discovery (insecure)
  - Authentication
  - Authorization
  - Key Management
- Authentication issues
  - Where does the cipher suite get negotiated?
    - Along with any other options
- What are the semantics for cipher suite negotiation?
  - If both support A and B, and prefer different algorithms, who wins - the initiator?
  - Discussion about how to structure the protocol
  - This is about infrastructure
  - There is some type of prioritization
  - The question is how to do the prioritization
  - To solve this problem can not consider as two separate conversations
  - If A talks to B and B is talking to A then must insure that it is a single conversation not two separate conversations
  - With control messages the issues is integrity not confidentially
  - There may be a confidentially issue with provider networks
  - Authorization will be tied to what you are willing to accept not to what you will transmit
  - The observation that integrity is related to what is received and confidentially is related to what is transmitted
  - There are cases where the receiver would be concerned with the received confidentially of data – credit card database
- Once authenticated a single time
  - Fast network recovery is a goal
  - Can the SA change?
  - The master key should be stored safely
  - What is the life time of fast reconnects?
    - This is key lifetime
  - Lots of discussion about the framework of fast reconnects
- Fast reconnects
- When shouldn’t a fast reconnect be used?
  - Boot strap an initial connection

    - Time to change the key
    - Key lifetime has limits based on amount of data and/or time
- Leave unspecified how the box gets the boot strap master key
- Central management is an issue
- We should not use EAP
    - No way to predict how often failures will occur
    - No way to determine if DoS
    - EAP will not get into the hardware
    - This way everything is a fast reconnect hence there is no need for higher layer entity
    - If higher layer things are required then it can do its thing
    - EAP was designed for dial-up to modem pool
    - Popular methods fail on shared media – prone to misuse
    - Customer interfacing versus infrastructure ports
    - No EAP methods support mutual authentication pre shared keys
    - Pass through model is not ideal
    - Does not support dual pass-through (switch to switch case)
    - We want to define a simple carrier mechanism so the upper level can use this mechanism
- AAA servers
- Towards a protocol
    - Many ways to do fast reconnect
        - Pick up the old connection where you left off
        - Use old key to create a new key and replace the old key
        - Use one key long term, just to generate transient keys
    - Third solution makes key management much easier
- Preliminaries for protocols
    - Master key is for a long term entity used to setup transient keys
    - Leverage GCM to provide secure transient keys
    - Number by use of nonce
        - Must handle nonce space used up
    - Reordering is the issue not loss
    - Need to consider the shared media case not just point to point case
    - Don’t want a group master
    - Master, transient key generator, transient key, and multicast key – 4 keys
    - Discussion of the keys
        - Master – manage authorization level (point to point)
        - Transient session key
        - Don’t use the master too much
        - But it is easy to change the master
        - KGK – key generating key
        - Not using security to create layer 3 functionality in layer 2 when stuff is discovered here then connect to it securely
        - 3 keys master, KGK, multicast
    - Back to protocol
        - How many KGK can be sent?

Why send more than one?

- Partial Protocol
  - Key id is key counter in the presentation
  - Discussion of the acknowledge phase
    - There are easy ways to handle the lost acknowledge – timeout or start receiving data
    - No data in a long time reset?
    - If link has been silent for some time what to do? Handle by policy not by protocol
  - Different use cases for authorization – what is a device allowed to do in an “unauthorized state” that is what does a bridge do with a BPDU when it has not been authorized. Edge devices probably don’t do anything. For the bridge probably in the form of you can not be root but you should listen to this control message.
  - Two levels of authorization – one is the level you get and another that you get after some level of integration. Can this continue that is can the level be increased? Yep should be able to.
  - Add some notion of registration number for cipher block
- How to set the Master Key
  - SNMP thing with heavy warning
  - Label on the box
  - Start with a weak key that has a low level of trust
  - Car wash code model for hot spots
- Issues
  - What happens if an attacker doesn’t allow B to respond?
    - A wants to avoid running out of nonces
    - Fall back challenge-response protocol
- Key management
  - Time to re-key
    - Independence of session keys
    - Does this need symmetric Diffie-Hellman?
    - Seeking time independence of the keys

Agenda for af for the rest of week

- Bob Moskowitz some thoughts on authentication, which may move to the discover portion
- Jim Burns has presentation for interface between AE and .1AF

Wednesday AM, May 19, 2004

Real Time Ethernet – Karl Weber and Ludwig Winkel

- Proposed Scope
  - Real-time communication for industrial automation applications
  - This standard extends 802.1D with optional short frames and scheduled transmission queue selection policies in a bridge for use in Industrial Automation applications requiring real time behaviors.
- Discussion
  - Time scale – what level for time sync. Below microsecond

Does the MAC have to be modified? Nope

What is the size of the network? Several hundred meters, typically 50 – 100 meters

Is 100 Mbps sufficient? For the foreseeable future

Would 802.3ah work? Shared bandwidth would be an issue and fiber optics can not handle the environment.

Observation – copper will have similar issues

Problem with the definition with real time – there is an issue with the definition of real time

- 802.1 should avoid having anything that looks like a real time standard

Observation – some components in the cited field bus standard cause concern – bit cut through and the topology is a ring

Is this doing the whole thing? Bit level cut through and ring redundancy

- Not every application needs this stuff
- Key issues is bounded time
- Other things depend on the application
- Power and chemical folks are looking for redundancy
  - There would be packet loss in some failover cases

802.1 does not want to do ring redundancy

There are several vendors doing a prototype deployment

- Want to standardize what has been proven in the field and standardize in the right place. If 802.1 does not want to do this then we will go find another place.

Issues with queue size and scheduling – if can not transmit normal traffic during a real time cycle then the normal traffic would exceed the current size of a queue in a bridge

Way to figure out way forward

- Suppose build a network that only will build real time traffic only
- This will really define the problem and adding the non real time traffic

What is the value of doing this work in 802.1?

- This is dealing with bridge like things
- Are there folks that will be here to do the work?
  - Need enough folks to do good and adequate reviews
  - In IEC it was not a problem but in automation only a few companies so it may be an issue for IEEE

If folks focus on the real time part and not the industrial control component then this project will spiral out of control so we must make sure that this stays focused on industrial control

This effort will succeed or fail based on the number of folks that do the work

If we were sure simple changing the queue scheduling would fix the problem then this would increase confidence that the work would be successful

To do this work it must be limited in scope and very focused

The next step

Organization structure – need a few that incredibly interested currently there are some that are mildly interested

Need total scope sorted out

ITU-T SG13 Liaison Letter about IEEE 802.1ab – Paul Congdon

Document is http://www.ieee802.org/1/files/public/docs2004/COM13-LS05.pdf And http://www.ieee802.org/1/files/public/docs2004/2-024Rev1.pdf

Requirements for network topology and resource status collection pertaining to the ITU-T draft new Recommendation (Y.12.qos) on a QoS architecture for Ethernet-based access networks

Have requirements on LLDP – additional TLVs, spanning tree state and port duplex setting

What they would like is a gateway that can access layer 2 managed information

LLDP is not a request response protocol and the collection of information is separate from the distribution

LLDP is point to point and does not forward across the bridge

SNMP does have its EtherType which could be use for this if the bridge supports it

There are established mechanisms to put together the topology of the network

There may be a bug in LLDP which is not getting the current link state for full and half duplex

LLDP allows organization to extend TLVs

Norm and Paul will work on a draft liaison letter to respond. At the July meeting we will create a formal response

The rules may allow Tony to respond as chair of 802.1 with “we discussed it and it will probably be a negative response.” Need a quick response so they are not waiting for us to give them a response.

Tony will respond as chair and we will make a formal reply in July but the response will most likely be no.

IEEE 802.1 af interfaces – Jim Burns

http://www.ieee802.org/1/files/public/docs2004/af_keyMgmt_IEEE_May2004_Barcelona.ppt

Assume MK to SAK process exists, how does .1af feed SAKs into the SecY?

Assume pre-shared MK exist, then don’t need authentication protocol

LMI Communication

- Modeled as shared data, indirect
  - Use get and set, modeled as data structure not as functions
  - Translates easily into state machine
  - Events occur based on setting data, actions then occur
- LMI from SecY to KaY
  - SecY reports the capabilities it supports
  - Which C.S., what is connectivity?
  - Encoding
    - Encoding for transmit
    - Encoding for enciphering
    - Why different?

        - Unique thing about this protocol (.AE), require to protect what's going out the door, so have to formulate the encoding, SECtag, before give to security engine to encrypt, so instead of being the last thing you do, encoding is the first thing you do.
        - Decide on value of variables, hand to encoding engine, hand to crypto engine
        - May have a delay, pipeline implementation
        - Next packet number to transmit on, the stuff in the crypto process or what’s being encoded, and may have been a key change in between the two
        - Need to explain the pipelining - encoding, enciphering, or going out on the wire
        - Need to get some things as a pair - the SA and the PN
        - Is there ever a case where the encipher SA IS different from the encoding SA?
    - Queue flush problem
    - Read from the KaY,
    - KaY doesn't need to know this level of detail
    - KaY only needs to know that the SA has changed and what new SAK is in use
    - SecY updates NextPN for each SA
- LMI from KaY to SecY
    - AN for each SA? Is it the SecY or the KaY that sets up the AN? Assume it's the KaY
    - SAK for each SA -- generate SAK whether valid or not, an invalid SAK should be a random number
    - If key management has fallen apart, still want to transmit frames,
    - For debugging, don't use zero as the number
    - Always fill in SAK field, if don't have one, put in a random one
- LMI from ??? to KaY
    - These are handled through management variables
    - Limit to number of allowed RX SCs, number of receive channels- a management variable, threshold which the implementation sets, but the manager may reset it. On a 10 Gbs link, it would be 1 RX channel.
    - Sliding scale of authorization, some ability of the box to say what level of authorization is required or desired.
- Chart for the LMI interface
    - Type of the variable, what it has to do with
    - Blue means transmit secure channel
    - Light blue is one SA within an SC, a SA
    - global data - yellow, CA level
    - Elements are management variables
    - PMK should be yellow not blue- a property of CA not SC
    - Receive SC is greenish
    - Discussion of PN
        - Transmit encoding for the SA

Transmit enciphering for the SA

Crypto on several frames at once, under key currently used, meanwhile getting frames from user. The PN relates to that, and not to the PN that you're currently encoding

TXSC Transmit SC, RXSC Receive SC

State of the SC - in the CA, necessary to keep state in transmit,

NotInCA -if receive, means get a pkt not in the CA. Having the SC is valuable for debugging. Can give valid debugging.

Maintaining old keys causes a security vulnerability.

Need some rules

When invalidate the KGK (key generating key), invalidate everything.

If MK becomes replaced, then invalidate

SC has ability to send a command -- Add SC to CA, remove from CA, stop using- can put the SC in the CA but not use yet, not till sure have symmetric connectivity

Need State diagrams for the interface, paths thru the state within an SC, store up to 4 SCs, on transmit side only need 2, the one your transmitting plus another.

on receive - use 3 at a time

On reception, nextPN is the next one going to receive

SA[0]

State is what state the key is in

Install means your calculating tables, etc.

Cmd can be sent to SA to install or uninstall key

Store 4 SAs

Start up

Just thinking this thru

new common port becomes available,

instantiate SecY and KaY on it

[changed during meeting - CA created with last saved value, may have gone down]

assume MK is per entity

have to be up in the 3 seconds it takes to reconfigure

could be the out of box MK

announcement occurs, [changed during mtg – no peer list - creates a peer list]

TX SC and RX SC created for each peer.

SCs created..

each key exchange results in SAK

when all peers have SC with SAK, and our TX SAK, know that peers are ready to receive our SAK

Events that Cause Action

New common port available

Empty peer list, or peer list could be out of date,

send announcement frame

SC with no matching peer in peer list

Don't be fast to remove from CA, device can come back up

Begin timer for removal

Install SAK

Bring up is optimistic - expect SC to be there, if not, then expect MK, if not then expect..

All peers, need to get set on receive side

use a little state machine for this

All peers… symmetry has been broken. A peer in list doesn’t have a SA, I’m transmitting... this is a problem

what to do? Uninstallkey.. provides a chain of actions

share, their MACOperStatus

Shared LAN with repeater in middle, when repeater on,

Requirement is that all the stations in one group see their MACOperStatus

unknown SA arrives

We will need a simulator

CA membership needs variables that may not be in SC

NotInCA has some of this

Sketch out required authorization for different clients

LAN-level Events

local station start

local station stop

peer enters CA

peer station leaves CA gracefully, requires a message –[no, later discussion]

peer station leaves CA ungracefully - if deal with this don't need to bother with graceful case

if CA becomes non-transitive or non-symmetric, then uninstall SA key for TX SA

MAC Operational set to false by SecY - no actions?

Choice of available Cipher Suite changes, disallowing the one I'm using

Questions slide

Whose job to ensure that symmetry and transitive attributes of CA are not violated? The KaY?

Which keys have lifetimes

SAK- PN wrap around, nothing else limits

MK - time, number of frames sent

If receiving SA approaches limit of PN should we attempt to initiate new SA creation. no. it's always the owner of the TX SA that creates a new SA.

How detect non-SecY neighbors?

KGK doesn't roll over in 4k years, only need one when master key changes, or if have timer..

We are master key acquisition method neutral- could be any D-H, Kerberos, etc.

Next steps

    - Further define variables needed
    - Develop SecY state machine
    - Define reference variables for LMI
    - Create state machines
    - Ensure all events needed for SecY are represented
- Outline of Doc
    - Protocol
    - Interface

Wednesday PM, May 19, 2004 Barcelona

Presentation, Jim Burns, Bob Moskowitz, Preeti Vinayakray-Jani

http://www.ieee802.org/1/files/public/docs2004/AFmay04Moskowitz-exchanges-v2.pdf

All the cases for getting MSK

- Pre-shared MSK case
- Point to point case
- System comes up no slower if there is someone without a pre-shared MSK than in the shared case
- I'm going to transmit using this KGK, I know about zero people on the LAN
- If you have the same key, validate it back
- Someone else responds
    - Send out again, see another person on the LAN
    - Continue till stop for some reason
    - In the point to point case, the count is 1
- Key exchange is sent as multicast
- What's the shared media?
- Should there be an initial announce message saying I'm up?
- First case, A and B know MSK
- Second case A knows MSK, but B does not
- Don’t do key exchange first, it's vulnerable
- John V. - easy if both parties have key
    - Easy if one has key
    - The only interesting case is, if both think they have the right key, but they are different
        - Someone changing key
        - Some sort of failure
- Might want to overlap MSKs.
- Key exchange and authentication machines are separate
    - When do you bring up the authentication machine?
    - John V. - the authentication model is not for us (IEEE) to specify
    - Mick- how often can we afford to do the authentication method? Depends on how cheap it is
    - We can give recommendations
    - Do the authentication method when don't have a key and need to have one.
- How pre-shared key gets there?
    - out of band?
    - in band?

trusted third party

Doesn't matter how the Master Key gets chosen, we say here's a transport mechanism

Authentication is expensive and subject to DoS

Could be via a layer 3 connection, an L3 protocol such as SNMP or RADIUS

Bob's presentation

Simple announce frame

Do we want L2 authentication? rather than L3?

At most we should provide a transport mechanism? Defining a transport for the authentication exchange

Mick – Assume transport is rate limited, only do X number of these types per second, say, 1 per second

Send key exchange protocol message

Put the authentication data at back of the field

If you get message and can't make sense of it, you pull authentication data, process key exchange, then look at authentication data

State machine- authentication comes in, have I done too many?

If not, deliver to authentication machine

In a transport, there's not an authentication

Authentication can come at any time

The only condition is whether got message from AE or not

We deliver a flag and nothing else

There is no proposal and response

Proposals and responses happen at L3

A way for 2 entities with L2 connection to not have to open an L3 connection

At end of key management, add authentication info

Key exchange can be run continuously

Steady low bandwidth channel running all the time

Key exchange protocol must be idempotent

Say we support 3 authentication mechanisms at L2

-the info transported,

-flag for state of current key,

-then sent to right place, standardized by someone else (?)

The important characteristic of carving it up this way is that correctness is a local property

All need to know about our machine is that it transports data at a rate not in excess of ..

Jim - this method means that you don't bring up L3 till you've gone through some security

All the attacks are at L3

L3 not accessible till AE is established

Authentication is a segregated process

Software is protected from a bad outside source getting into L3 stack, without going through security minded software

Minimizing your attack surface

The paper details these steps

We will need standard management protocols to deploy this stuff

Back to SecY Management Parameters

Thursday AM, May 20, 2004

AF protocol Continued - Jim Burns

http://www.ieee802.org/1/files/public/docs2004/af_keyMgmt_IEEE_May2004_Barcelona.ppt

The announce must contain the KGK and number of receivers, actually the key exchange is the announcement

LAN should be ready in 4 messages but there cannot be a state

It is not possible to have a starting state and then a transfer state because you can not know when to transition from start state to transfer state.

Also, it is not possible to know when to start because the network will always be in same state of flux.

What is the cost of processing malicious incoming keys? Not much, it is a look up so there is not some DoS attack.

Lots of discussion about how to balance protocol needs with security needs

The interesting case is when both sides think they have the correct key but the keys are different. Otherwise, things are rather straightforward

Should have the data around to answer the question how did I get here?

When should authentication start?

Depends upon how expensive it is to start and run. We can make recommendation and start when we know that we do not have a key.

.1af trigger events for state machines - Bob Moskowitz

This is the first pass at trigger events

“MAC up” – you can transmit

Need a “begin” signal that can reset all of the state machines

Do we believe there is a benefit to layer 2 authentication?

Assume the transport is rate limited. Put the authentication data in the key exchange

One constraint can send the same thing

Layer 3 is not accessible until .1af has done its thing

This discussion needs to be captured in 802.1af

1af trigger events for state machines – Bob Moskowitz

http://www.ieee802.org/1/files/public/docs2004/AFmay04Moskowitz-exchanges-v2.pdf

This is the first pass at what trigger events

MAC up – you can transmit

Need a begin signal that can reset the all of the state machines

Do we believe there is a benefit to layer 2 authentication?

Assume the transport is rate limited. Put the authentication data in the key exchange

One constraint can send the same thing

Layer 3 is not accessible until af has done its thing
This discussion needs to be captured in 802.1af

Attendees:
Paul Bottorff
Jim Burns
Paul Congdon
Kevin Daines
Sharam Davari
Arjan de Heer
Thomas Dineen
Anush Elangovan
Hesham Elbakoury
David Elie-Dit-Cosaque
Jee Sook Enn
Maria Esteve Lloret
Norm Finn
Yukihiro Fujimoto
Steve Haddock
Tony Jeffree
Tetsuya Kawakami
Loren Larsen
Yannick Le Goff
Bill McIntosh
Katsuya Minami
Dinesh Mohan
Bob Moskowitz
Satoshi Obara
Don O'Connor
Hiroshi Ohta
Glenn Parsons
Allyn Romanow
Dan Romascanu
Jessy V Rouyer
Ali Sajassi
Dolors Sala
Mick Seaman
Alon Shavit
Yoshihiro Suzuki
Geoff Thompson
Michel Thorsen
Genadi Velev
John Viega
Preeti Vinayakray-Jani
Karl Weber
Ludwig Winkel

May 2004 Barcelona, Spain

Michael D. Wright