

Zonal Address Partitioning in the Local Space

Document Number: IEEE 802-ec-14-0071-01-00EC

Date Submitted:

2014-11-03

Source:

Roger B. Marks

EthAirNet Associates

404 Montview Blvd

Denver, CO 80207 USA

Voice:

+1 802 capable

E-mail:

roger@ethair.net

*<http://standards.ieee.org/faqs/affiliationFAQ.html>>

Re: Proposed P802c PAR

Venue:

Discussion session on 802c, use of the "locally administered address space" and privacy

2014-11-03, San Antonio, TX, USA

<http://www.ieee802.org/secmail/msg17996.html>

<http://www.ieee802.org/secmail/msg18019.html>

Purpose:

To elucidate issues related to the proposed P802c PAR, to clarify the implications of Company ID assignments by the IEEE Registration Authority, to review the factors related to random address assignment, to demonstrate the value of innovation in using the local address space, to demonstrate the need for standardization to enable innovation in using the local address space, and to propose zonal addressing.

Notice:

This document represents the views of the author and is offered as a basis for discussion.

Zonal Address Partitioning in the Local Space

Roger B. Marks
EthAirNet Associates

Table of Contents

1.Executive Summary

2.Details

Key Points

- Proposed P802c PAR proposes partitioning of the local space to incorporate IEEE CID-based address allocation.
- Some 802 participants advocate a flat 46-bit space to maximize use with random assignment.
- This proposal suggests that:
 - partitioning has many valuable uses (exemplified here by zonal addressing and switching)
 - P802c PAR should enable use of local addresses, not merely allocation.
 - Truly enabling use will require IEEE 802 standards.
 - IEEE 802 protocol should use clean blocks, not isolated “company IDs”.
 - Partitioned space in the core can coexist with flat space at the edge and is compatible with fully randomized use.

CID

- IEEE RA defines CID to include a 22-bit space.
- IEEE RA has allocated public CIDs in the “10” quadrant of that 22-bit space.
- The first 14 CID assignments preclude clean block assignments in the “10” quadrant.
- IEEE 802 protocol should use clean blocks.
- Coordination with the RA needed to ensure that CIDs are not assigned in space used by IEEE 802 protocol.

CID Assignments

Local Quadrant						
00	xxxx	00 10	xxxx	xxxx	xxxx	xxxx
01	xxxx	01 10	xxxx	xxxx	xxxx	xxxx
10*	xxxx	10 10	xxxx	xxxx	xxxx	xxxx
11	xxxx	11 10	xxxx	xxxx	xxxx	xxxx

*Current public CID assignments are in this quadrant, and widely dispersed in this quadrant.


A Flood becomes a Torrent?

- Internet of Things (IoT) device access might be rare, so the destination address may frequently be unknown; this could lead to a high proportion of frame flooding.
- Privacy advocates are suggesting randomized addresses and frequent address changes. This would result in many frames with unknown destination addresses, and therefore to flooding.
- We need to plan to avoid a torrent of flooding.

Zonal switching instead of address learning

Core Switch


Address is partitioned to incorporate a zone identifier.
Provides topology-based switching.
Switch selects a port based on zone identifier in the zonal DA.
Switch learns zones, not addresses, based on zonal SA.
Much smaller forwarding tables: one entry per zone.
Flooding results from an unknown zone (rare), not an unknown DA.




zonal switch

Access Point (AP)


Zone 1
conventional bridging locally



Zone 2
conventional bridging locally



Zone 3
conventional bridging locally



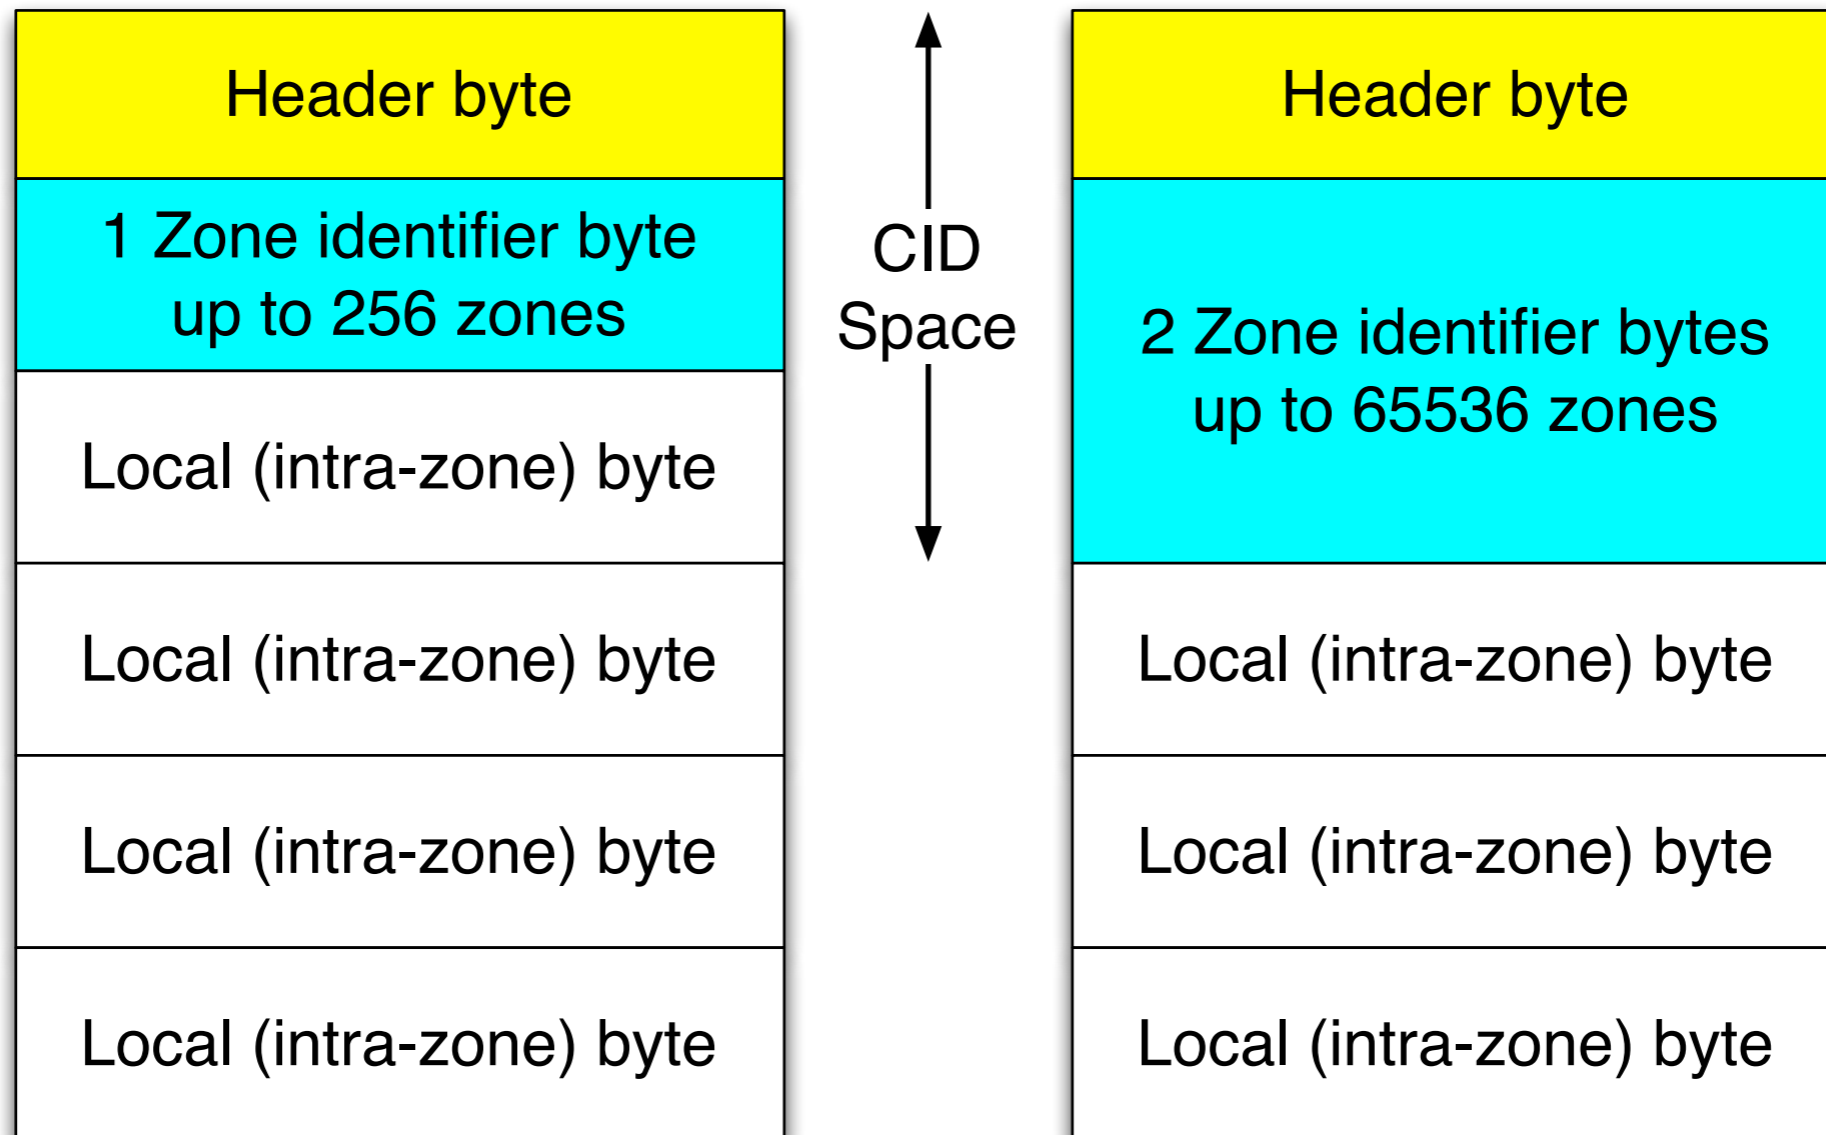
User

User User User

User User User

User User User

Example address structures for zonal subnetting




Header could also be a half-byte, etc.

Example: One-Byte Header

Bytes of Zone ID (if zonal)	structure identifier	quadrant identifier	Local bit	Multicast bit
00=1 byte 01=2 bytes 10=3 bytes 11=4 bytes	00=zonal 01=reserved 10 reserved 11 reserved	01 or 11 if RAC agrees to avoid CID allocations (10 is already off- limits, and 00 is slightly problematic)	1	0


Address Aliasing

Core Switch

Address is partitioned to incorporate a zone identifier.
 Switch selects a port based on zone identifier in the zonal DA.
 Switch learns zones, not addresses, based on zonal SA. 

← Core address

Access Point (AP)

Local mapping table maps each internal (user-side) address to a corresponding external (core-side) alias address.
 Access point replaces access address with core alias address in transmissions to core, and vice versa. 

← Access address

User


terminal selects among many (46?) random bits
 or could use a global address

Core address	Access address
alias unique within the core	local (could be terminal-selected) or global
alias unique within the core	local (could be terminal-selected) or global

A kind of address translation but, unlike NAT, it's one-to-one.


Address Updating

Core Switch

Address is partitioned to incorporate a zone identifier.
 Switch selects a port based on zone identifier in the zonal DA.
 Switch learns zones, not addresses, based on zonal SA. 

← Core address

Access Point (AP)

Local mapping table maps each internal (user-side) address to a corresponding external (core-side) alias address.
 Access point replaces access address with core alias address in transmissions to core, and vice versa. 

← Access address

Core address	Access address
alias unique within the core	local (could be terminal-selected) or global
alias unique within the core	Unique, once assigned by AP; could be identical to Core address

User

terminal selects among many (46?) random bits or could use a global address	later, AP assigns unique address (could be identical to Core address)
--	--

Conclusions

- Standardization of the local address space structure offers rich opportunities to develop enhanced Layer 2 networking; e.g., topology-based zonal addressing and switching, that is unavailable with the existing global address space.
- Structuring the local address space need not harm the use of that space for access or inhibit privacy improvements and is compatible with 46-bit random addressing for access, using methods like addressing aliasing.
- The “10” quadrant of the CID is fine for address allocation using individual RA-assigned CIDs but appears to not be a suitable basis for IEEE 802 to build local address structures for IEEE 802 protocols.
- IEEE 802 protocols should use of a clean block of addresses outside the 10 quadrant, coordinating with the RAC so that CIDs are not allocated there.

Recommendations

- We should applaud the efforts made by the RAC to encourage exploitation of local space.
- Standardization of the local address space structure should proceed, with 802 working hand-in-hand with the RAC to promote success.
- The P802c project should be open to both CID and non-CID solutions. CID-based solutions should consider the space outside the “10” quadrant of the CID, and coordination with the RAC should be initiated in preparation to request assignment of a clean block.
- PAR should be revised to more explicitly support the use as well as the allocation of local addresses.
- Local address space structure standardization should not be considered in any way “mandatory,” since local usage should be up to the local administrator. However, standardization will play an enabling role in applications such as new switching methods and new random assignment mechanisms.

Proposals - Draft P802c PAR

- 5.2.b: This will allocate a portion of the address space for protocols using an IEEE Registration Authority assigned Company ID. Another portion of the local address space will be allocated for assignment by local administrators. **A portion will be allocated for use by IEEE 802 protocols using a partitioned local address.**
- 5.5: This project will enable protocols that automatically configure **and use** addresses from a portion of the local address space.
- CSD should be revised to include address usage.

Details

- Background materials and details are provided on the following slides.

Scenario

- Half of IEEE 802 address space (with the 7th bit set on) is locally administered and not presumed globally unique.
- The IEEE-SA Registration Authority (RA), managed by the RA Committee (RAC), is assigning 24-bit CIDs with the 7th bit set on, and suggesting that these could be used to form 802 addresses.
- A proposed P802c PAR would initiate a project to provide recommendations and rules for the 802 local address space, allocating a portion for protocols using the CID and another portion for local administration.
- Local space is being considered for temporary address assignments to alleviate privacy concerns related to static global addresses based on EUI.

Divergence of Views

- Some P802c proponents suggest dividing the local space using the CID, perhaps as an identifier of an address assignment or resolution protocol.
- Some, considering privacy, wireless, and access issues, prefer a wide-open, unstructured space to maximize entropy and privacy while minimizing collision probability.
- The two sides have been standing in opposition.
 - Few signs of convergence of views.
- Should be possible to reach a common understanding, considering:
 - people are addressing different problems
 - the local space is *local* and does not require a universal solution
- Can a resolution of the conflict offer lead to better use of the local space?

What is a CID?

- a bit confusing; RA documentation does not fully explain
 - <https://standards.ieee.org/develop/regauth/tut/eui.pdf>
- RA specifies that a CID is 24 bits, with the last two bits of the first byte set to 1 and 0
 - if one were to build an 802 MAC address using the CID as the first three bytes, that MAC address would have the local bit on and the multicast bit off
- But people are talking about a “quadrant” of the CID space. What’s that?

Public IEEE RA CID Allocations

as of 2014-11-03

0000	1010	0101	0111	0010	0100
0010	1010	1111	1101	0110	1010
0110	1010	0100	1001	0111	1101
0110	1010	1110	0110	0100	1010
0111	1010	0100	1011	1010	0100
1001	1010	0011	1000	0100	1011
1001	1010	1101	0010	0011	1111
1010	1010	0000	1111	1000	0101
1100	1010	0010	1100	0110	1100
1101	1010	1010	0001	0001	0011
1101	1010	1111	1111	0110	1011
1110	1010	0010	0010	1011	0100
1111	1010	0100	1100	0100	1101
1111	1010	0011	0111	0110	1111

Observations:

- First two LSB in first byte are 802-compatible (local bit set, multicast not).
- Next two LSB in first byte are all “10”; RA seems to staying in this quadrant.
- Structured use of any other bit leads to collision since RA has allocated every option.
 - RA has even allocated every possible PAIR of bits in the first byte.
- *Possible inference*: The RA has staked out every corner of this “10” quadrant and does not want any trespassing! Bits will not be designated. Only solo addresses will be assigned. Don’t expect a clean block assignment in this quadrant.

“Local” OUIs

- As [noted by the RA](#), “A very small number of assignments made prior to adoption of IEEE 802 standards have the X bit equal to one (BlockID assignments). These assignments are documented in the CID registry.”
- However, they do not appear to be included in the CID registry.
- The list of 18 can be found using a grep developed by Mark Rison:

```
grep '^ .[2367abABefEF]\- ' oui.txt|cut -b6-10|while read f; do echo ===== $f; grep '^ ..-$f' oui.txt; done
```
- 5 are in the “10” quadrant.
- 13 are in the “00” quadrant.

0000	0010	0000	0111	0000	0001
0000	0010	0001	1100	0111	1100
0000	0010	0110	0000	1000	0110
0000	0010	0110	0000	1000	1100
0000	0010	0111	0000	0000	0001
0000	0010	0111	0000	1011	0000
0000	0010	0111	0000	1011	0011
0000	0010	1001	1101	1000	1110
0000	0010	1010	1010	0011	1100
0000	0010	1011	1011	0000	0001
0000	0010	1100	0000	1000	1100
0000	0010	1100	1111	0001	1100
0000	0010	1110	0110	1101	0011
1010	1010	0000	0000	0000	0000
1010	1010	0000	0000	0000	0010
1010	1010	0000	0000	0000	0001
1010	1010	0000	0000	0000	0011
1010	1010	0000	0000	0000	0100

What is a CID again?

- It's important to differentiate the CID space (22 bits) from the currently-allocated "10" CID quadrant (20 bits).
- The current draft P802c PAR would "allocate a portion of the address space for protocols using an IEEE Registration Authority assigned Company ID."
- The CID is not limited to the "10" quadrant.
- Utilizing the other portions of the CID space will require early RAC coordination.

P802c in the Data Center

- Example: Fiber Channel over Ethernet
- Possible usage:
 - CID allocated for a protocol (not to a “company”)
 - address server allocates addresses built from assigned CID
 - Other users sharing the local space will avoid assignments using that CID and therefore avoid collision
- consult experts for the correct explanation

Internet of Things (IoT)

- The current P802c draft PAR says that “Such protocols will allow virtual machines and IoT devices to obtain a local address without local administration.”
- It is not clear that the IoT address acquisition problem is parallel to that of the data center.
- IoT would represent large numbers of devices at the edge of a network.
- We generally expect a large number of IoT devices.
 - this may overload switches, as well as overloading address space and address acquisition methods
- How can we help?

A Flood becomes a Torrent?

- IoT device access might be rare, so the destination address may frequently be unknown; this could lead to a high proportion of frame flooding.
- Privacy advocates are suggesting randomized addresses and frequent address changes. This would result in many frames with unknown destination addresses, and therefore to flooding.
- We need to plan to avoid a torrent of flooding.

Sample WLAN Environment

- Some privacy advocates have argued for the maximum number of available bits (i.e., 46, since two bits of the 802 address space are otherwise designated) to minimize collision possibility in a single large switching domain, particularly when local addresses are assigned randomly. Example of scale:
 - [Santa Clara CA football stadium wireless LAN:](#)
 - 68,500 seats
 - 1,200 access points
 - 1 access point per 57 seats

Collision-Free Probability for Random Assignment

A = number of addresses available

U = number of users

1 address assigned randomly per user

Collision-Free Probability: $P(U,A)$

$$P(U,A) = \frac{A!}{(A-U)!A^U} \approx \exp\left(\frac{-U \cdot (U-1)}{2A}\right)$$

Source: E. W. Weisstein, *CRC Concise Encyclopedia of Mathematics*, Second Edition

Collision in Stadium Scenario

- terminal selects from among 2^{46} available addresses
- terminal associates with an AP
- address collision with any other terminal is possible over the air (not only within the core), even if the other terminal associates with a different AP, because frame over the air may be heard anywhere
- special problem for shared medium

What else can we do with partitioning?

- Partitioning of the 46 bits available in the local space is advocated as means of differentiating address acquisition methods by protocol (e.g., in data center).
- If we are going to partition, let's see what else it's good for.
- Let's consider partitioning local addresses by topology.

Bridging gets messy in a large zone

Single zone

i bits to work with

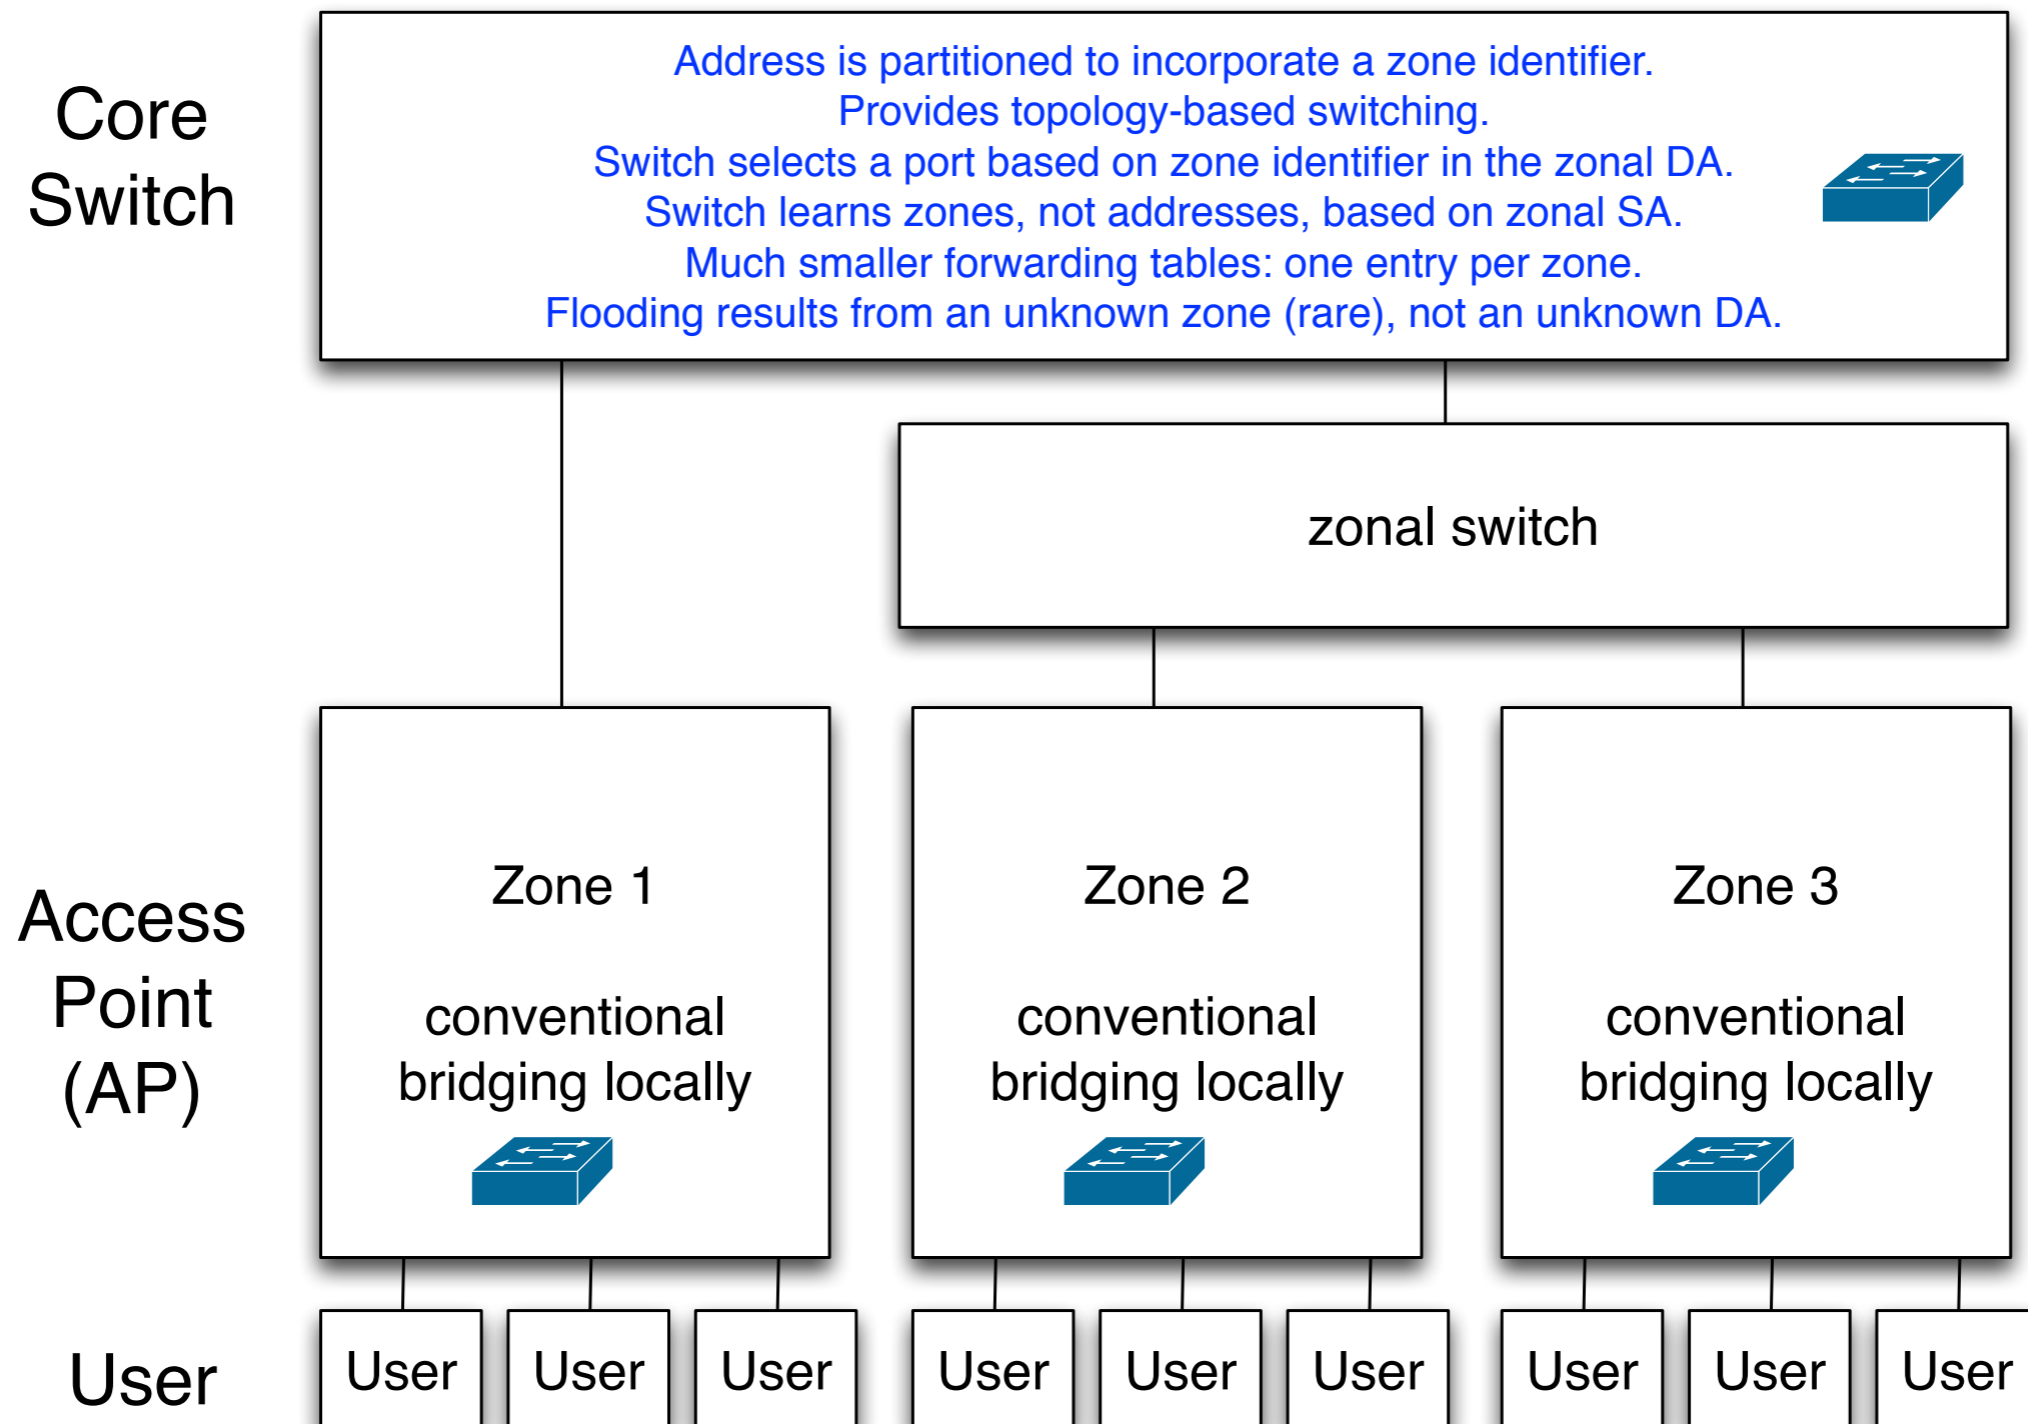
$d=2^i$ addresses



Switch uses bridge learning and stores (e.g.) $2^{16} = 65,536$ addresses in filtering database.




Zonal switching with a topology-based hierarchical address instead of address learning



Some advantages of zonal addresses

- Switches need learn only the association between the port and the zone, instead of the association between the port and the address.
- the zone is less costly to store, since it is shorter than the entire address.
- the switch needs to learn only one port association for each zone, instead of one per address, so forwarding tables are much shorter and less frequently updated
- far less flooding, because an unstored address need not be flooded as long as the zone is known
- Note:
 - the zone ID has significance only within the switching domain
 - all the zones in the switching domain need unique IDs

Example address structures for zonal subnetting



Header could also be a half-byte, etc.

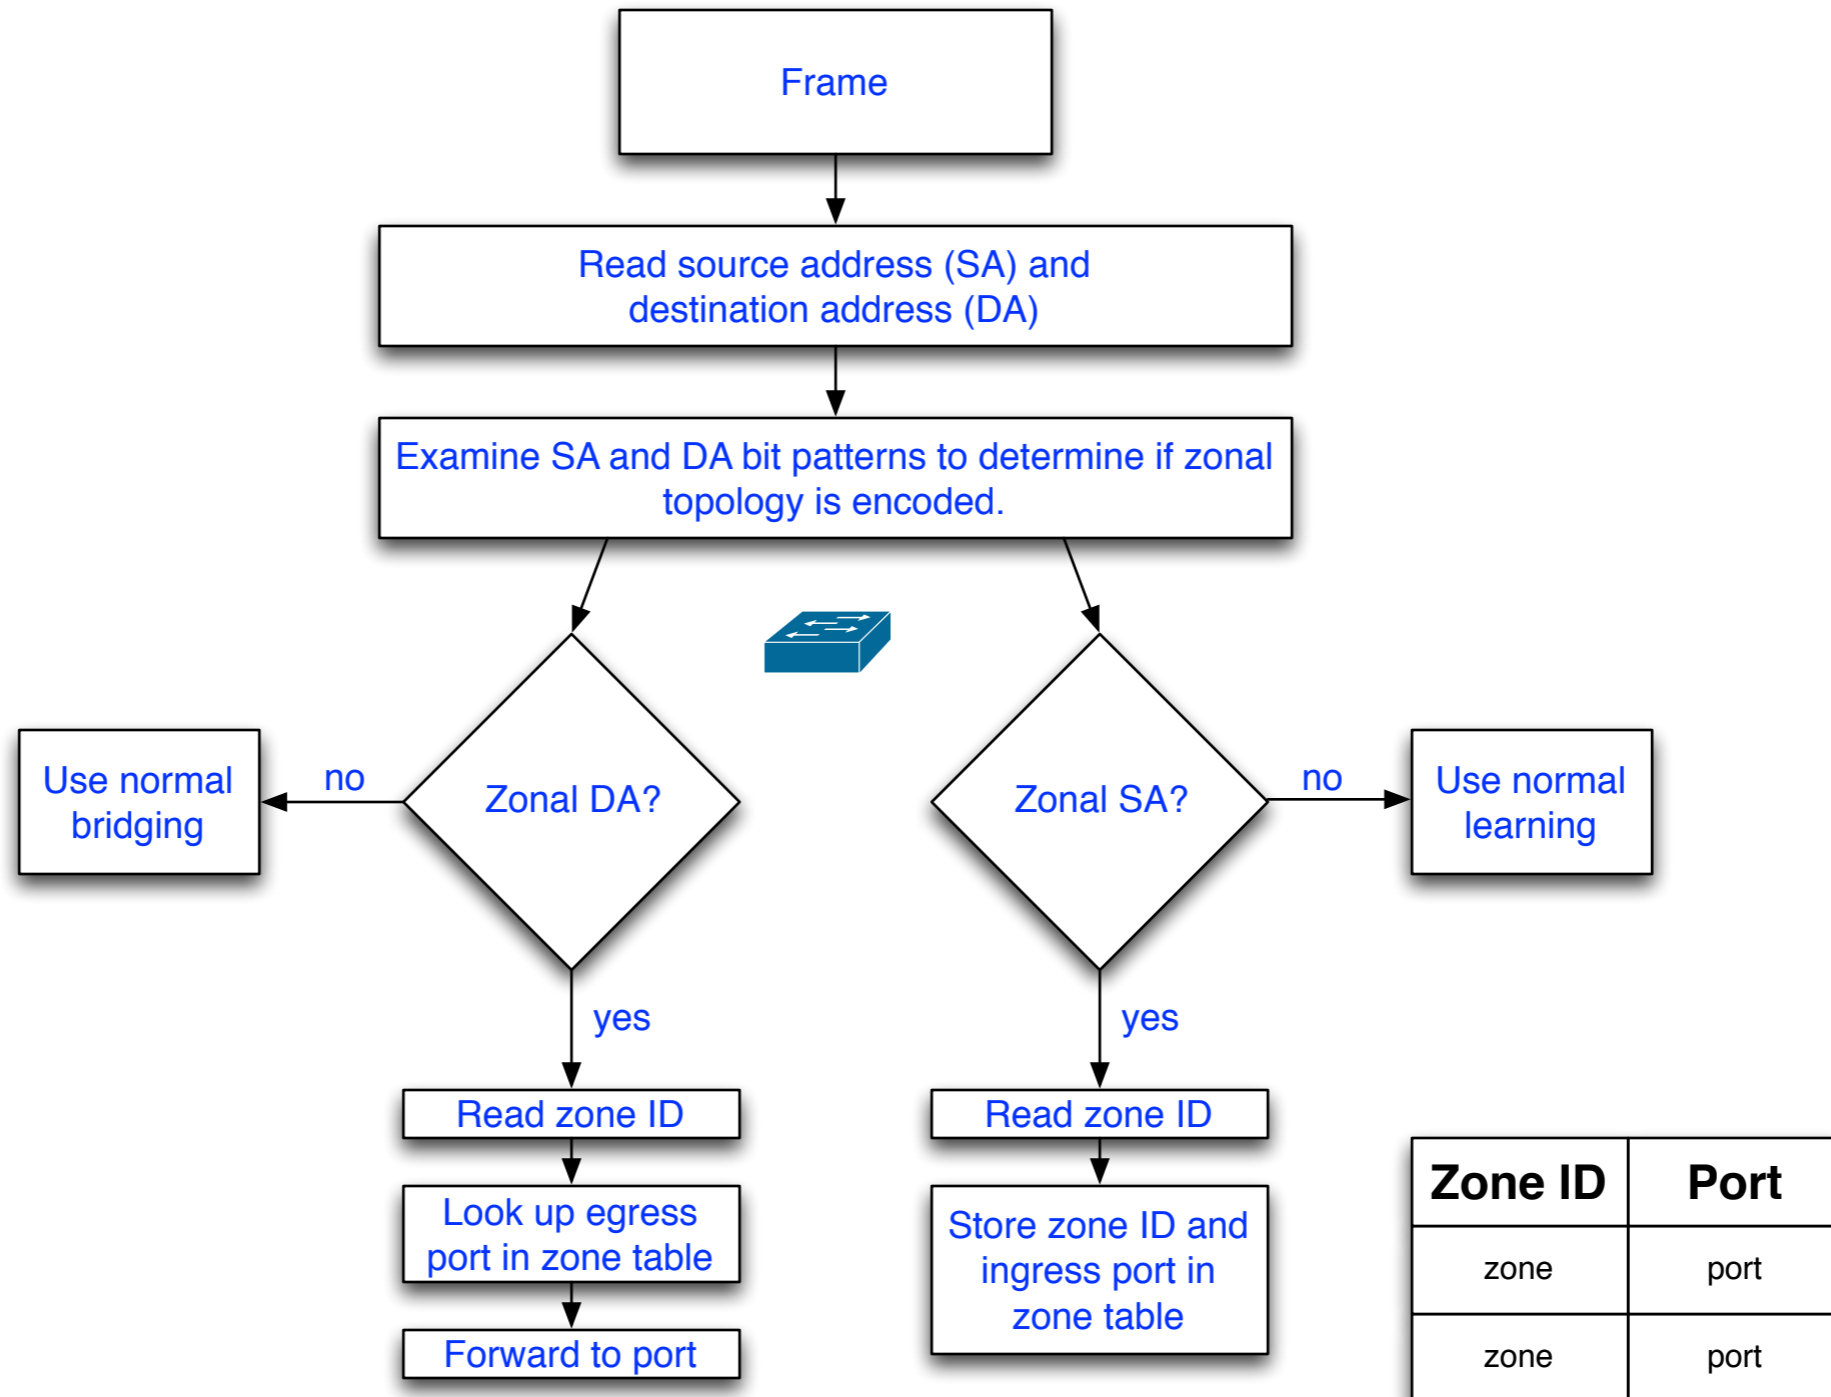
Requires a standard partitioning plan

- In order to support zonal subnetting, the switch needs to ascertain whether the address is zonal
 - and if so, parse the address

Example: One-Byte Header

Bytes of Zone ID (if zonal)	structure identifier	quadrant identifier	Local bit	Multicast bit
00=1 byte 01=2 bytes 10=3 bytes 11=4 bytes	00=zonal 01=reserved 10 reserved 11 reserved	01 or 11 if RAC agrees to avoid CID allocations (10 is already off- limits, and 00 is slightly problematic)	1	0

Switching based on mixed conventional and zonal addresses



RAC Coordination

- Use of the header in conjunction with an RA-assigned CID would require coordination with the RAC.
- The suggestion would require, in effect, a block of CIDs, not just a CID.
 - e.g., a block of CIDs two or four bits wide
 - to be used in header as well as zonal and local addresses
- So far, the RA has been assigning CIDs in the 10 quadrant of the CID space.
 - The 10 quadrant could not be used without affecting current assignments.

Reminder of why “10” quadrant is Off Limits

Public IEEE RA CID Allocations (as of 2014-10-20)

0000	1010	0101	0111	0010	0100
0010	1010	1111	1101	0110	1010
0110	1010	0100	1001	0111	1101
0110	1010	1110	0110	0100	1010
0111	1010	0100	1011	1010	0100
1001	1010	0011	1000	0100	1011
1001	1010	1101	0010	0011	1111
1101	1010	1010	0001	0001	0011
1101	1010	1111	1111	0110	1011
1110	1010	0010	0010	1011	0100
1111	1010	0100	1100	0100	1101
1111	1010	0011	0111	0110	1111

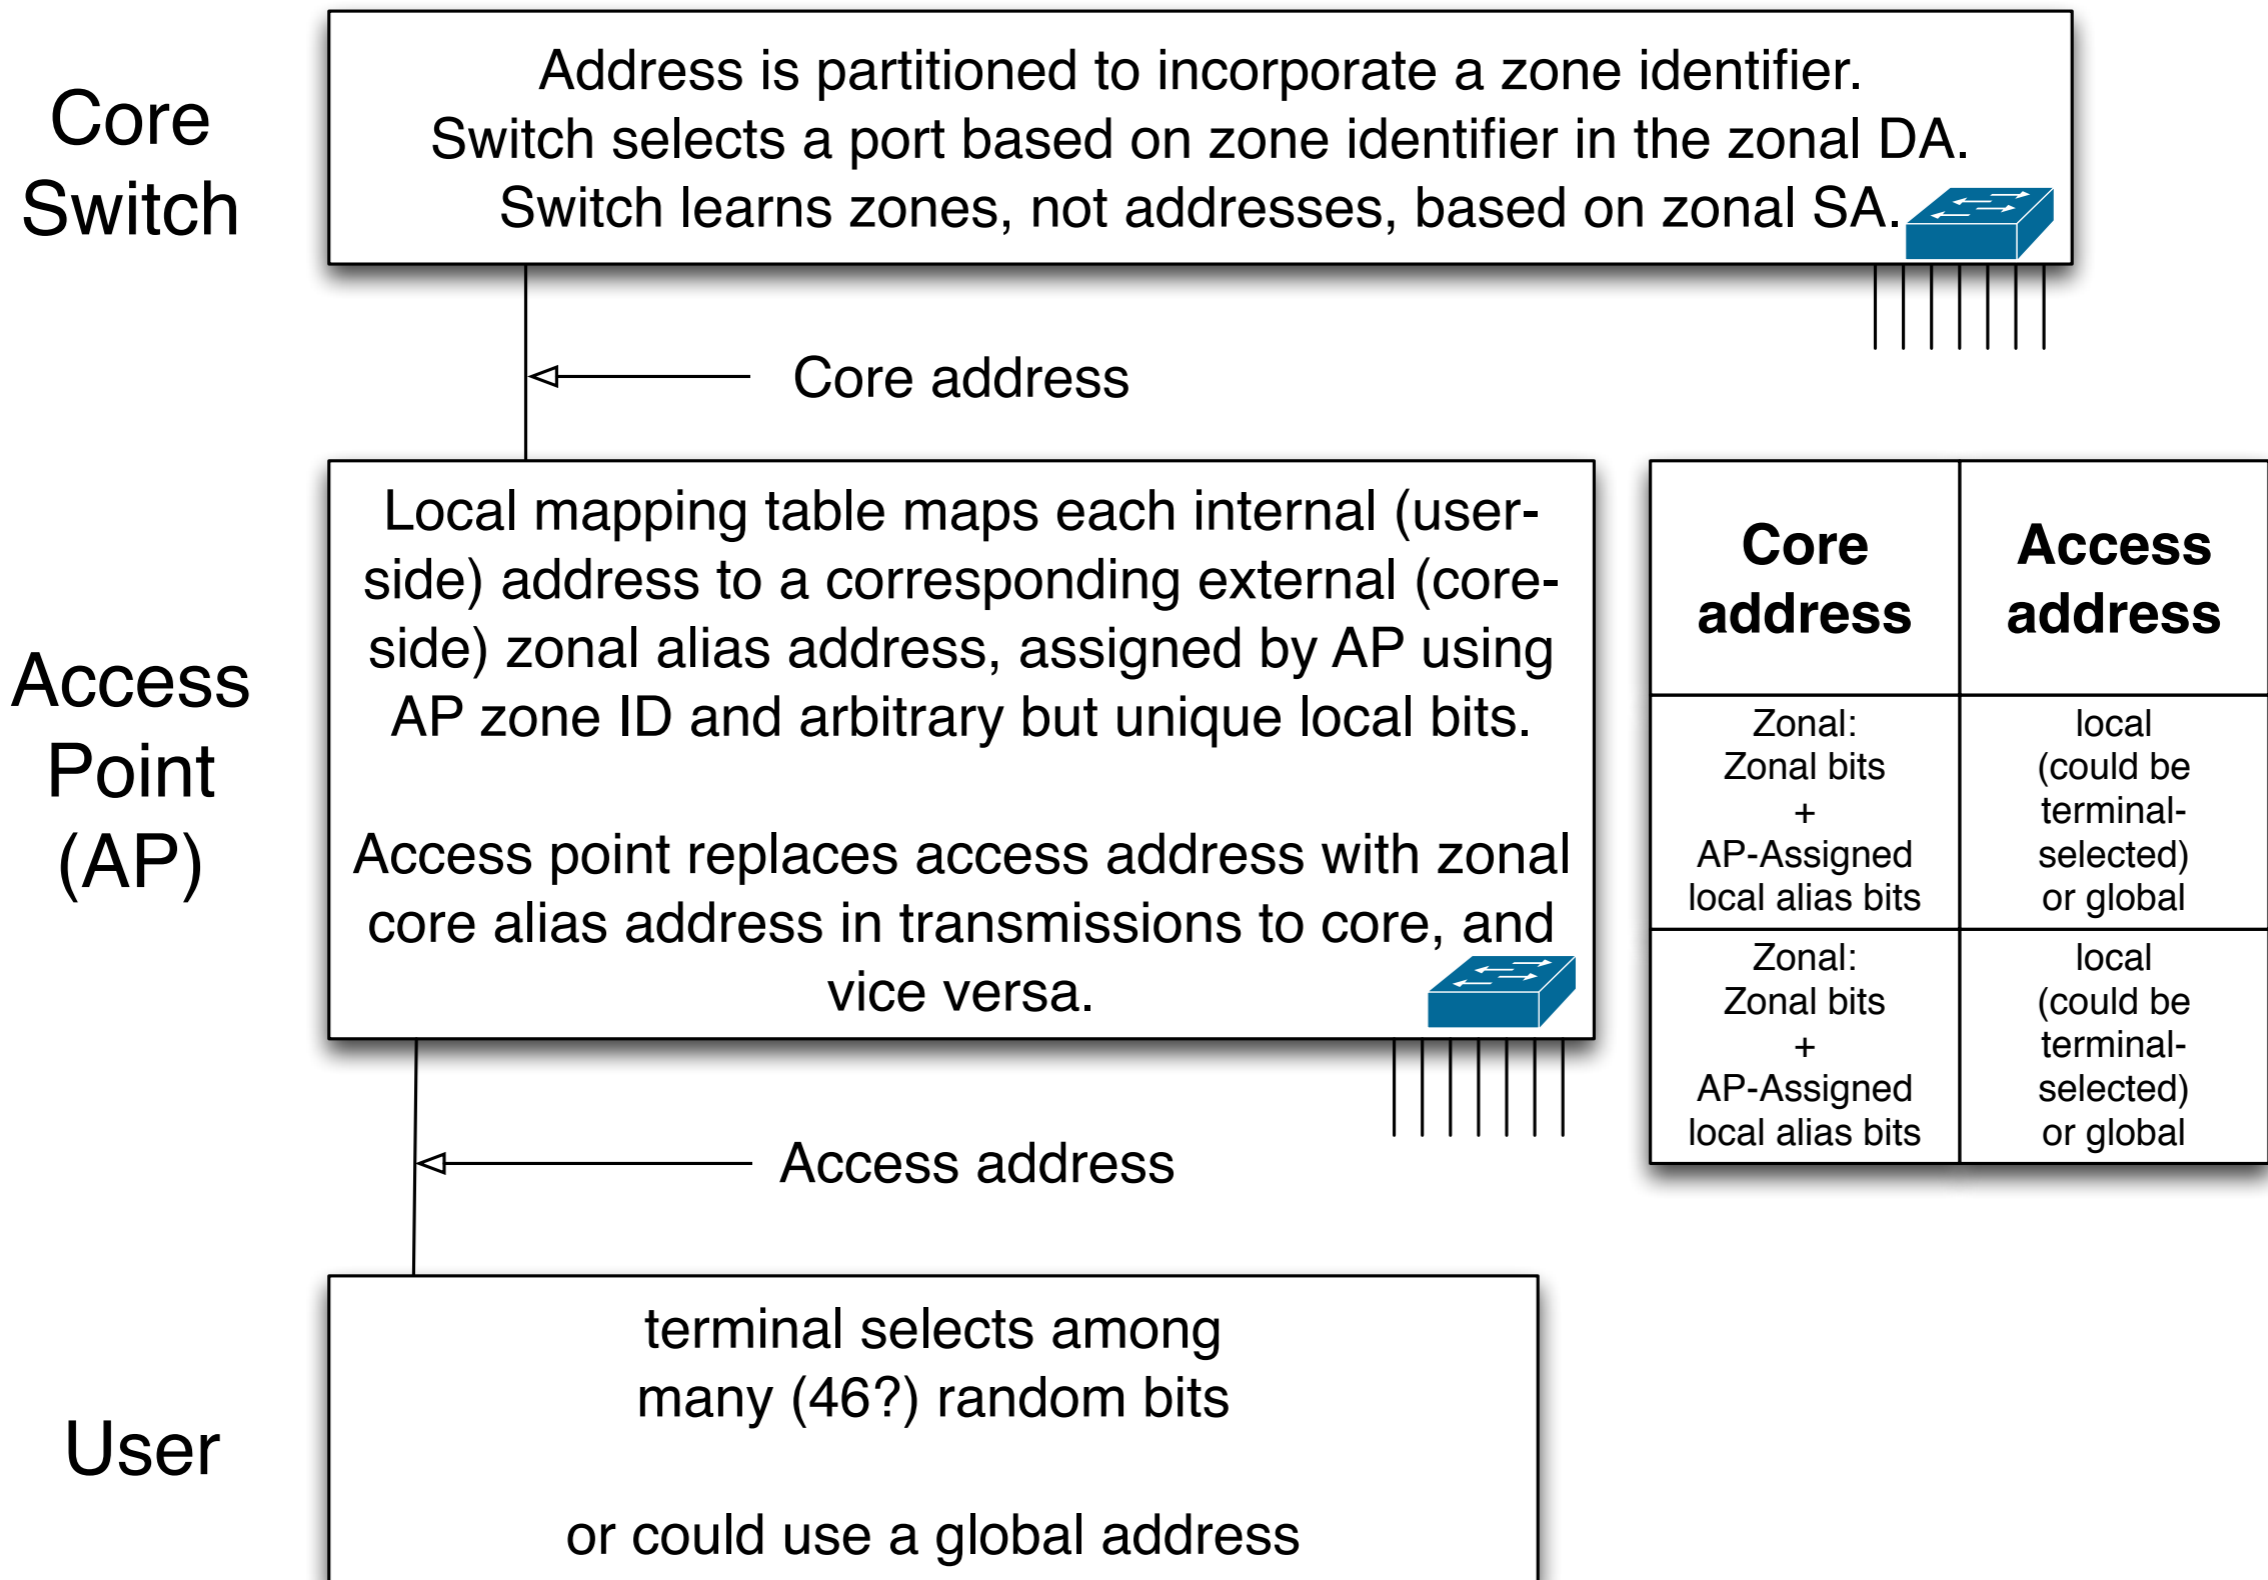
Conclusion:

- Any structuring should be done in another quadrant, or else structuring should be achieved without using the CID.

Is this consistent with random assignment?

- Some have proposed that devices randomly select a local address from the 46 bits not already specified by 802.
- Is a structured local address consistent?
- Would the structure limit the size of the random space and lead to unacceptable collision rate?
- How would the device know how much of the space is allocated to zonal addresses?

Zonal Address Aliasing



Example: Stadium WLAN

- $2^{16} = 65,536$ users
- $2^{10} = 1024$ APs
- 1 AP per $2^6 = 64$ users
- 1 zone per AP
- $2^{10} = 1024$ zones
- 2^{16} zone IDs (enough to support 65,536 zones)
- 2^{24} local IDs (each AP can support 16 M users)


Improving Collision-Free Probability in Stadium Scenario

- If each terminal choose one address, collision-free probability is:

$$P(U, A)=P(2^{16}, 2^{46})=\frac{2^{46}!}{(2^{46}-2^{16})! \cdot (2^{46})^{2^{16}}} \approx 0.999969$$

- Becomes worse if devices change addresses.
- To improve the odds, use the random address as only a temporary one.
- Could provide a protocol for the AP to assign a guaranteed-unique zonal address, such as, for example, the Core address.
 - Send a message to the terminal, using the temporary address, ordering a change to the assigned zonal address. Terminal acknowledges, and AP updates database.

Zonal Address Aliasing, rev



Conclusions

- Standardization of the local address space structure offers rich opportunities to develop enhanced Layer 2 networking; e.g., topology-based zonal addressing and switching, that is unavailable with the existing unstructured global address space.
- Structuring the local address space need not harm the use of that space for access or inhibit privacy improvements and is compatible with 46-bit random addressing for access, using methods like addressing aliasing.
- The “10” quadrant of the CID appears to not be a suitable basis for IEEE 802 to build local address structures.
- If 802 is to make use of the CID, it should seek an allocation of a clean block of addresses outside the 10 quadrant.

Recommendations

- We should applaud the efforts made by the RAC to encourage exploitation of local space.
- Standardization of the local address space structure should proceed, with 802 working hand-in-hand with the RAC to promote success.
- The P802c project should be open to both CID and non-CID solutions. CID-based solutions should consider the space outside the “10” quadrant of the CID, and coordination with the RAC should be initiated in preparation to request assignment of a clean block.
- Local address space structure standardization should not be considered in any way “mandatory,” since local usage should be up to the local administrator. However, standardization will play an enabling role in applications such as new switching methods and new random assignment mechanisms.