DOCSIS Service Identifier (SID) Space

John J. Downey – Sr. CMTS Technical Leader

March 10, 2020

Agenda

- SID Background
- SID Exhaustion
- SID Recovery Suggestions
- Implementation Plans
- Non-Immediate Solutions
- Multiple MAC Domains Pros and Cons
- Customer Example

SID Background

CISCO SID Space is Per MAC Domain

- 13 bit field = 0 – 8191
  - ✓ 0 used during CM registration
  - ✓ Few on high-end reserved leaving ~ **8175** total
    - ➢ IE: 8191 used for FFT for US Spectrum viewing
- Spec says 8K "**must**" and 16K "**may**"
  - ✓ Gray area of spec doesn't say per domain or per US
  - ✓ Others may have 16K Expanded SID space
- Every US Service Flow ID (SFID) has unique SID
  - ✓ **Note**: DS SFs not the issue so Multicast does not count in that limit
- Previous discussions to make per US or keep per domain and increase to 16K, but limited demand and too much work
- Command to find total:
  - ✓ `show interfaces Cable 1/0/0 sid | count enable`

DOCSIS MULPI 7.2.1.2 Spec

- Five types of Service IDs defined
  1. 0x3FFF - Broadcast, intended for all stations
  2. 0x3E00-0x3FFE - Multicast, purpose is defined administratively, refer to Annex A
  3. 0x2000-0x3DFF - Expanded Unicast, intended for particular CM or particular service within that CM, when supported by both CM & CMTS
  4. 0x0001-0x1FFF - Unicast, intended for particular CM or particular service within that CM
  5. 0x0000 - Null Address, addressed to no station
- CMTS **MAY** support Expanded Unicast SID space
  - ✓CMs **MUST** support
- Unicast SIDs (including Expanded Unicast SIDs) assigned by CMTS **MUST** be unique on given logical US
  - ✓CMTS **MAY** support Unicast SID assignments which are not unique within single MAC-sublayer domain as long as they are unique on given logical US

SID Exhaustion

SID Exhaustion Factors

- DOCSIS Settop Gateway (DSG) devices increase SIDs
  - ✓Usually D2.0 CMs embedded and very low speed requirements
  - ✓Every house could have 3 CMs (2 DSG and 1 HSD CM)
- US bonding with "`sid cluster 2`" exacerbates issue
  - ✓Especially with low speed flows for maintenance & other services
- "Stale"/unused service flows
- Offline CMs hold onto 1 SID each
  - ✓Offline for 24 hrs unless `ccm off delete` used or linecard failover
- Future Factors
  - ✓Gaming with Low Latency DOCSIS (LLD) in future using Proactive Grant Service (PGS) will be another US SID
  - ✓Mobile (5G) Backhaul

Stale Dynamic Service Flows

- Dynamic SFs may not tear-down according to T8 timer creating “stale” service flows
- `show int cable 5/0/0 service-flow | inc S\(d\) US.*d`
- Will show dynamic service flows that have been up for days
- The `\` is needed to include `(` and the `.*` is wildcard till “`d`” seen
- **Note**: Needs to be done on all DS interfaces

SID Cluster Background

- `cab sid-cluster-gr num-of-clus`
  - ✓Default is `1` & `dynamic`
- SIDs assigned during registration of CM
- 2 SIDs assigned if Max US Rate >=28 Mbps
  - ✓Max Information Rate (MIR) as configured in cm file
- 1 SID otherwise
- Best practice suggests `2` for better per-CM US speed
  - ✓`cab sid-cluster-gr num-of-clus 2`
  - ✓Other option is increase cm file US Max Traffic Burst to 30 kB or more
  - ✓**Note**: D3.0 uses Continuous Concatenations and Fragmentation (CCF), which means no need to change Max Concat field from default of 3044, but 8k suggested in case D3.0 CM registers in D2.0 US mode
- **Potential Issue** – all BE US service flows that are mtc-mode (US bonded) will use 2 SIDs for each service flow

Cisco Implementation of SID Cluster

- Current SID Usages (SW Arch Model)
  - ✓Primary SID is synonymous with CM; Primary SID = CM
  - ✓SIDs are used for SF identification
  - ✓SID is used for quick look-up for both control and data paths features
  - ✓MAC/IP host chain and entry for IPv4 and IPv6 are intimately tied to SID value
- Reuse SID within SID Cluster to conserve SID space
  - ✓Assign same SID for all US chs of SID Cluster on bonded flow
  - ✓For example, SF1 has two SID clusters and US BG with chs 1 & 2

	SID Cluster 0	SID Cluster 1
US ch 1	SID = 1	SID = 2
US ch 2	SID = 1	SID = 2

- `cbr8#sh cable modem 10.12.1.106 service-flow ver | in SC`

```
      Sid Cluster    : SC-0, Sid [ 1 1 1 1 1 ]
      Sid Cluster    : SC-1, Sid [ 12 12 12 12 12 ]
```

SID Recovery Suggestions

Potential Solutions – Recover Unused SIDs

- Proactively “clean up” “stale” and Offline service flows
  - ✓ Scripting
- Add CMTS global config so flows with no activity > 300 sec are torn down if CM/eMTA does not do it automatically
  - ✓ `cable service flow activity-timeout 300`
- Manual process to clean up dynamic SFs with no traffic (US & DS)
  - ✓ Needs to be done on all DS interfaces
  - ✓ `sh int c5/0/0 serv 1321 count verbose`
    - ➢ Used to verify no bit rate before tearing down
  - ✓ `test cable dsd “mac address” “service flow”`
    - ➢ Used for dynamic service flow deletion
- Linecard failover will remove `offline` CMs as will `ccm off del`

SID Utilization Reduction Opportunities

- Assuming all segmentation has been exhausted
  - ✓Only 1 node per SG
- Revert back to default SID Cluster of `dynamic`
  - ✓Increase Max US Traffic Burst on high-speed tiers
    - ➢ Service class names (SCN) for cm file easier to implement
    - ➢ **Note**: CSCvs98821/Aha! CBR8-I-221 filed & implemented in 16.12.1z
      - ❖ User-Configurable Dynamic SID Cluster Threshold per Mac Domain applies to Max US Rate threshold when 2 SIDs assigned
- Review need for 4 Mbps maintenance service flow
  - ✓Add 1-2 Mbps US to everyone and get rid of this SF
- Review the need for nRTPS call signaling flow

Example of 7141 SIDs; (7141/8175 = 87%)

Name	MD	CPE	Total SFs	NRTPS	Other	Off	BE	NB	UB	SIDs	>28M
CTS	7/0/2	2100	4150	379	3	6	3762	837	2925	7141	2

- US bonded SFs will use 2 SIDs
  - ✓ 2925*2 = 5850 SIDs + 379 nRTPS SFIDs + 837 NB SFIDs (probably DSG devices) + 6 offline = 7072 – not quite 7141
    - ➢ Could have some dynamic SIDs for UGS (VoIP) happening also
- Potential Savings = UB SFIDs minus service flows > 28 Mbps
  - ✓ SFIDs < 28 Mbps would go from 2 SIDs down to 1
  - ✓ 2925-2 = **2923** SIDs potentially saved if we go to dynamic SID config
  - ✓ 7141-2923 = 4218/8175 = 51.6%
- Even if half those get gaming SFs later, still have plenty
  - ✓ 7141-(2923/2)/8175 = 69%

cBR-8 Scaling

- 64,000 per chassis & 10,000 per LC; 8,000 for LCHA
- 1500 per MD due to ranging
  - ✓ Not communicated & maybe not an issue with so many primary DSs now
- Service Flow Scale
  - ✓ SUP-160 supports 400,000 DS Service Flows
  - ✓ SUP-250?
  - ✓ 100,000 DS Service Flows are available for every 2 LC slots
  - ✓ 20,000 DS Multicast Flows per cBR
  - ✓ No US SF limit from SUP
  - ✓ Linecard supports 128,000 US + DS SFs
- Classifier Scale
  - ✓ No practical system or LC limit
  - ✓ Tested 256,000 Classifiers on system
- BSoD L2VPN
  - ✓ 32,000 supported (MPLS)
- Routing Table
  - ✓ No fixed limit
  - ✓ 1,000,000 IPv4 and IPv6 routes supported in any production config

Implementation Plans

Immediate Plans – Dynamic SID Cluster

- Identify cable interfaces with extreme SID amount (>5700)
  - ✓ `show inter cx/y/z sid | count enable`
- Identify total UB flows on those interfaces
  - ✓ Should save that many minus handful user flows > 28 Mbps
- Do speed tests on high-end tiers to verify proper speed
  - ✓ If not, identify, rectify or make note
- Document wideband CM count
  - ✓ `sh cable modem cx/y/z wideband | count w`
- Remove and apply Dynamic SID Cluster on specific interfaces
  - ✓ `no cable sid-cluster-group`, then `cable sid-cluster-group dynamic`
- Reboot those wideband CMs
  - ✓ `clear cable modem cx/y/z wide reinit`
- Verify SID space dropped to expected
- Do speed tests on high-end tiers to verify proper speed
  - ✓ If not on previous good tested CMs; identify, verify proper US bonding mode, then increase Max US Traffic Burst

Operational Considerations

- Future Plans and Architectures
  - ✓Stay proactive on SID space recovery & limits for future SF plans
  - ✓Home gateway solution & maybe no DSG
  - ✓DAA with R-PHY will be smaller SGs
    - ➢Each SG usually single mac domain
    - ➢**Note**: 204 MHz US and D3.1 OFDMA will require DAA
  - ✓Cloud architecture will not have this issue
    - ➢Virtual MAC vs physical MAC
- Double mac domains in each SG
  - ✓Can do now, but complex & not Keep it Simple Solution (KISS)!
- Cisco looking into doubling SID space to 16K
  - ✓Under investigation, affect on existing CPE unknown

Multiple Mac Domains
Pros and Cons

Cons of One Mac Domain per SG

- 16 USs allowed per domain
  - ✓ Issue with architectures of 1:4 and more than 4 USs per fiber node (FN) needed
  - ✓ Max of 12 SC-QAMs and 4 OFDMA per US Controller
- More USs in node = more UCDs & DS Map traffic
  - ✓ Typical .4 Mbps of DS overhead for each US
    - ➢ OFDMA US & LLD 1 ms Map idea will be worse
  - ✓ More UCDs can create longer registration
- 8175 SID limit per domain shared across all USs

Pros of Two Mac Domains per SG

- Supported today
- Bonding across domains not an issue
- More SID space supported
  - ✓SIDs needed for DSG, D3.0 SID cluster 2, LLD PGS, UGS
- Less DS overhead since less USs & UCDs per domain
- Supports 1:4 architecture w/ 6 SC-QAM+2 OFDMA USs per FN
- Assuming node is 1:2 US segmentable
  - ✓If not, then 2 domains will require node split or more US spectrum
- Staggering DSs between domains achieves better distribution;
  - ✓D2.0 load balance
  - ✓D3.0 BG selection & load balance
  - ✓Maybe resilient bonding groups (RBGs)
- **Note**: Be careful with Restricted LBGs for D3.0 CMs
  - ✓All DSs for bonding must be in RLBG

Cons of Two Mac Domains per SG

- Less efficient D1.x & D2.0 LB
  - ✓ Modems in certain FNs can't use all DS freqs
    - ➢ Maybe less of an issue with D2.0 attrition
- D3.0 primary distribution uneven (staggering DSs may help)
  - ✓ Resilient Bonding Group (RBG) affect?
- DSG box registration may act erratic
- More cable interfaces
  - ✓ cBR-8 supports 16 total cable interfaces (SGs); 32 planned for R-PHY
    - ➢ Matches US connectors/FNs
- UCDs from 2 domains may increase CM registration
  - ✓ DS freq override (DFO) should help
- May not be supported for CLI Simplification feature
- NB CMs (D1.x & 2.0) may get in bad cycle if primary goes down & DFO keeps pointing back to bad DS

Customer Example

Situation and Recommendation

- **Problem**
  - SID utilization beyond 70% threshold on 24 MAC domains in specific region
  - All on Cisco cBR-8 (5999 city1 + 5664 city2)
  - Accounts for 0.206% of all MAC domains on cBR-8s in this region
- **Contributing Factors**
  - SID maximum limit per MAC domain is 8175
  - SID Cluster 2 (originally recommended by Cisco)
    - ✓ Uses 1 additional SID for customers not in top tier plan and 1 additional SID for maintenance BE flow
  - Stale SIDs
- **Recommended Short Term Actions**
  - Config Dynamic SIDs – Immediately available, reduce SID utilization by ~30-40%
- **Results by Leveraging Dynamic SIDs without Splitting Nodes**
  - 30-40% reduction in SID utilization per MAC domain
  - Worst case - SIDs reduced from 7585 to 4608, 41% reduction in SIDs
- **Next Steps**
  - Feature for user-defined US threshold to dynamically add second SID
  - Customer input on feasibility of Dynamic SID configuration
  - Technical collaboration on long term management

Sample SID Usage & Potential Savings

- CM with 5 SIDs

SID	Mac Address	IP Address	Type	US SFID
2013	dead.beef.cafe	10.100.117.2	BE	2871
8138	dead.beef.cafe	10.100.117.2	BE	2871
8139	dead.beef.cafe	10.100.117.2	BE	11245
8140	dead.beef.cafe	10.100.117.2	BE	11245
8141	dead.beef.cafe	10.100.117.2	NRTPS	11246

  - ✓ SFID 2871 uses BE at 4 Mbps for US Maintenance
  - ✓ SFID 11245 uses BE at 10 Mbps for customer US high speed data
  - ✓ SFID 11246 uses nRTPS for call signaling

- After Dynamic SID applied and CM rebooted

SID	Mac Address	IP Address	Type	US SFID
655	dead.beef.cafe	10.100.117.2	NRTPS	12544
1730	dead.beef.cafe	10.100.117.2	BE	12543
2013	dead.beef.cafe	10.100.117.2	BE	2871

Customer Dynamic SID Cluster Implications

- Qualification and decision to go with Dynamic SIDs
- Decision to go footprint-wide or targeted area
- Wideband CMs need to be reinitialized when Dynamic SIDs configured
- Estimated effort to configure Dynamic SIDs for top 24 MAC domains; ~ 2-3 hour maintenance window
- Monitoring top tier subscribers to ensure US speed is being achieved

Closing Thoughts

- Expanded SID Space
  - Currently unknown how CMs and DSG devices will react to Expanded SID space
- SF growth will be from:
  - Upgrades (2.0 to 3.0 to 3.1)
  - Service tier increases
  - Adding more US services with associated flows
  - Device count in existing SGs
- **Note**: Some D3.0 STBs present as well with 3 BE flows using 6 SIDs!