Apricot 2012, New Delhi, India

Route Convergence Monitoring & Diagnostics

Clarence Filsfils (cf@cisco.com)

Ketan Talaulikar (ketant@cisco.com)

February 29th, 2012

Agenda

- Routing Convergence
- RCMD Overview
- Routing Convergence Measurements
- Application
- Reporting For ISIS
- A Case Study
- Insights into Routing Convergence

Routing Convergence

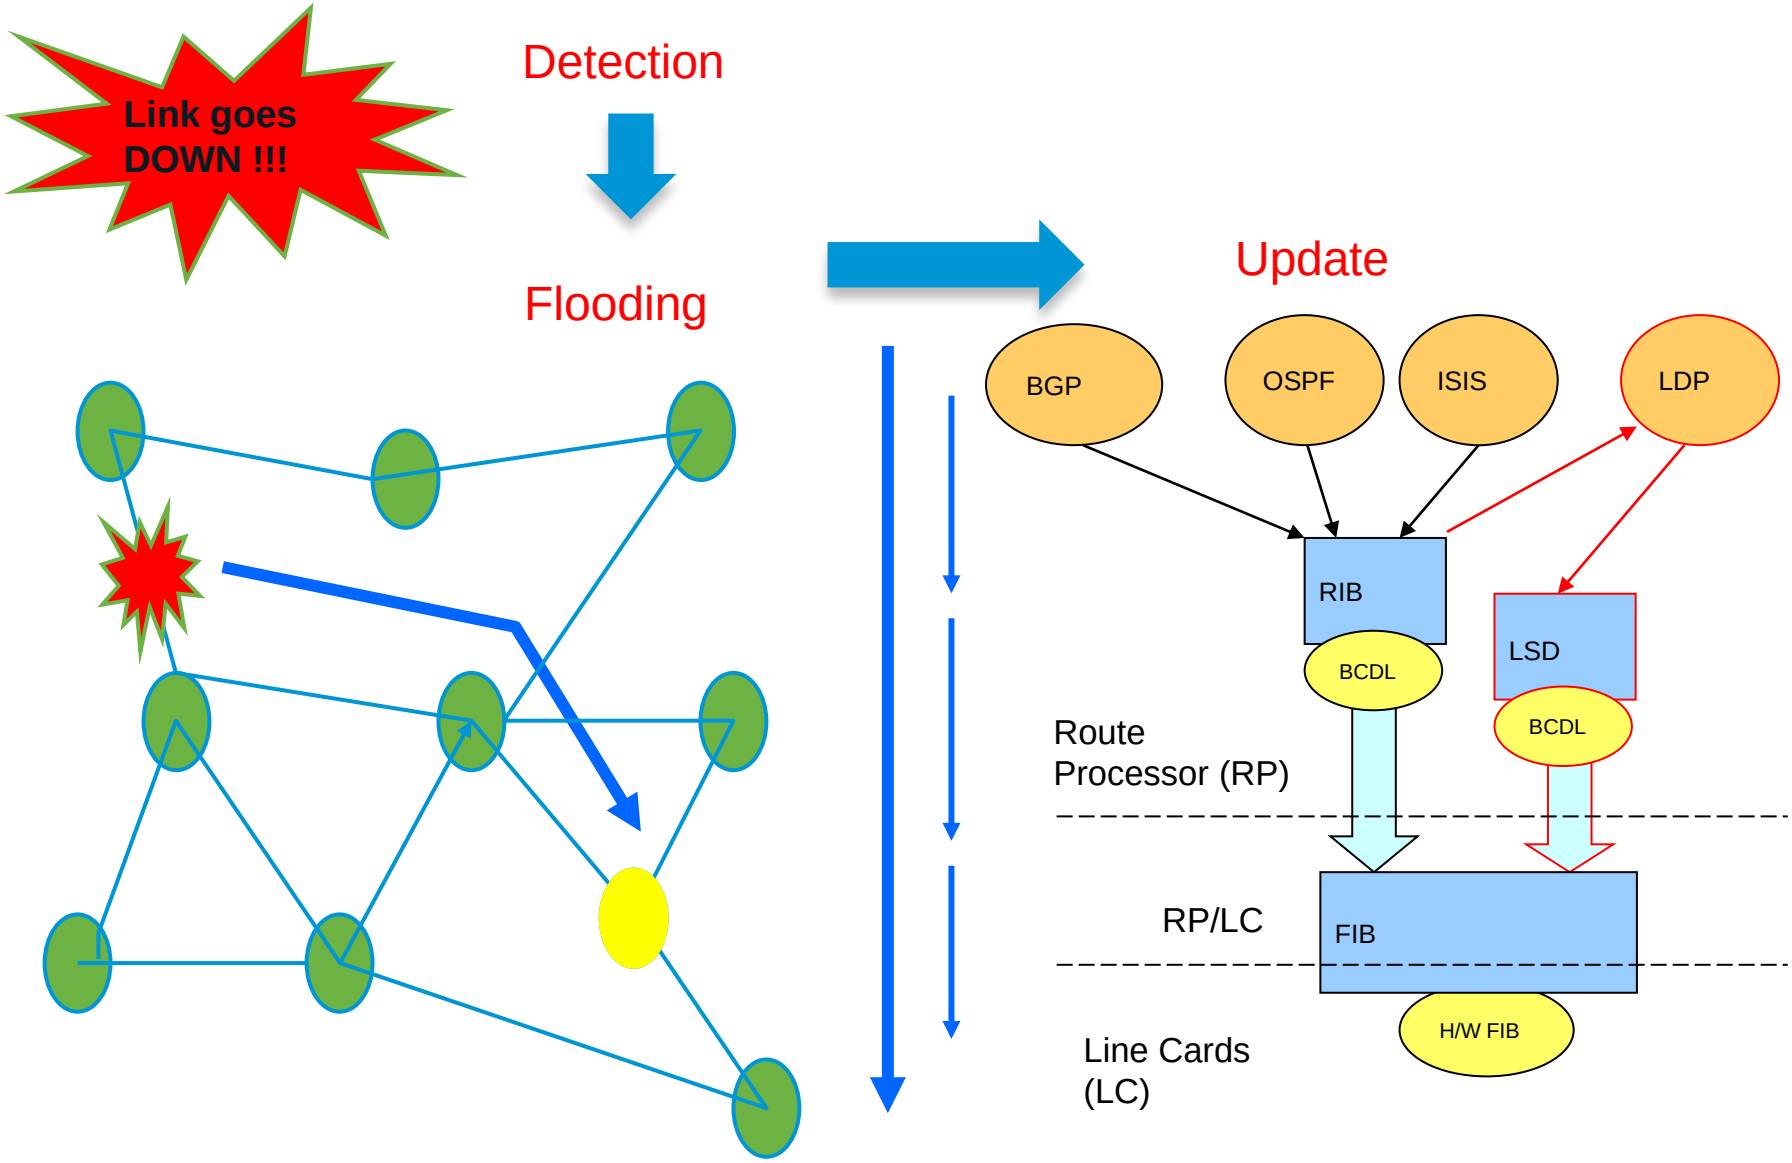

Routing Convergence – Questions …

- Was end-to-end connectivity restored within a second?
- What is the network availability for the last N days?
- How do network design changes affect convergence?
- What is the route change flooding propagation delays seen in the network?
- Are timers and other tuning parameters working optimally?
- How are different routers or segments of network handling failures?

How can we get these answers in production networks?

Routing → Monitoring → Measurement → Analysis → Diagnostics

RCMD Overview

- A tool that collects and reports data related to Routing Convergence
  - Provides "in-router" view of convergence events – data exported via XML can be correlated & leveraged by an "offline" tool
  - Lightweight and always-on – non convergence impacting
  - Persistent – archived for use after hours/days
  - Covers SP core IGP/LDP network in first phase
- Runs in two modes
  - Monitoring – detecting events & measuring convergence
  - Diagnostics – additional (debug) info collection for "abnormal" events
- Debugging router/network-wide convergence for an event is complex
  - Affects ISPs ability to commit to SLA

Flooding Measurement

- Flooding propagation delays across routers
  - Timestamps of Type 1-2 LSA (or LSP) change detection
  - Approximate Process/Rx time of LSA/LSP in OSPF/ISIS
- For each LSA/LSP that was flooded
  - determine the origination time at failure point
  - for each remote router, determine the flooding time
  - flooding = Time until the remote router got the LSA/LSP
  - case flooding >= 100msec: orange flag
  - case flooding >= 200msec: red flag
  - compute average and percentiles…

Update Measurement

- Correlation data for tracking cause of SPF events across routers
  - Determine Type 1-2 LSA (or LSP) regenerated by router(s) connected to failure
  - Determine SPF events in which these LSA/LSP were processed on each router
- For each router, for each SPF event
  - Check duration until all Critical prefixes were updated across all line-cards
  - Check number of Critical prefixes that were affected (say “C”)
    - case C > 1k: “scale is larger than expected”
    - case C <=1k & update >= 250msec: orange flag
    - case C <=1k & update >= 500msec: red flag
  - Compute average and percentiles…

Application: Sub-second Convergence?

- If
  - detection is known to be < 10 msec
  - flooding for any IGP event was verified by RCMD to be < 200 msec
  - update for any IGP event was verified by RCMD to be < 500msec
- Then
  - for any IGP event, for any router, loss of connectivity < sec
- This does not require any complex offline processing

Application: Fine-grained Analysis

- 3rd Party Network Monitoring Tools could retrieve RCMD data to compute exactly when each router finished processing each event and hence determine
  - when exactly the connectivity was restored (likely much less than the sub-second bound)
  - whether there were loops during the IGP convergence
  - further analysis …
- This is possible – “offline”

ISIS* Convergence Measurement

- Monitoring intra/inter-level & external routes on a per SPF basis
  - tracking on prefix priority basis
  - maximum of 4 sets of routes tracked per SPF– one for each priority
- Covers all types of SPF
  - Full
  - Incremental
  - Partial Route Calculation
  - Nexthop Change Calculation
- Reporting done on per SPF event basis
  - aggregate convergence time also reported on prefix priority basis
  - convergence time for a priority is when last route (intra/inter/ext) that is provisioned
- Provides timers values applied for the SPF along with activity statistics and trigger reasons & times

* Also support OSPF

RCMD Report for ISIS* SPF Event

- ISIS convergence events (i.e. SPF runs) and time taken to provision route+label changes across all LCs
- SPF computations statistics, trigger reasons, wait times
- LSPs that were processed and the timestamps of when their change was detected
- Route prioritization aware – reporting done on aggregate route priority set and not per prefix
- Leaf network deletes detected during the SPF (throttled)
- Statistics – route counts, LSP change counts

* Also support OSPF

A Case Study – Flooding

```
-------------------------------------------------------------------------------------------------
 Event Num  |       LSP ID        |Seq Num |SPF Run | Lvl |        Time         |    Trigger
-------------------------------------------------------------------------------------------------
      13      0020.0203.2002.00-00     50       0      L2   Feb 16 14:39:55.802       aj if
```

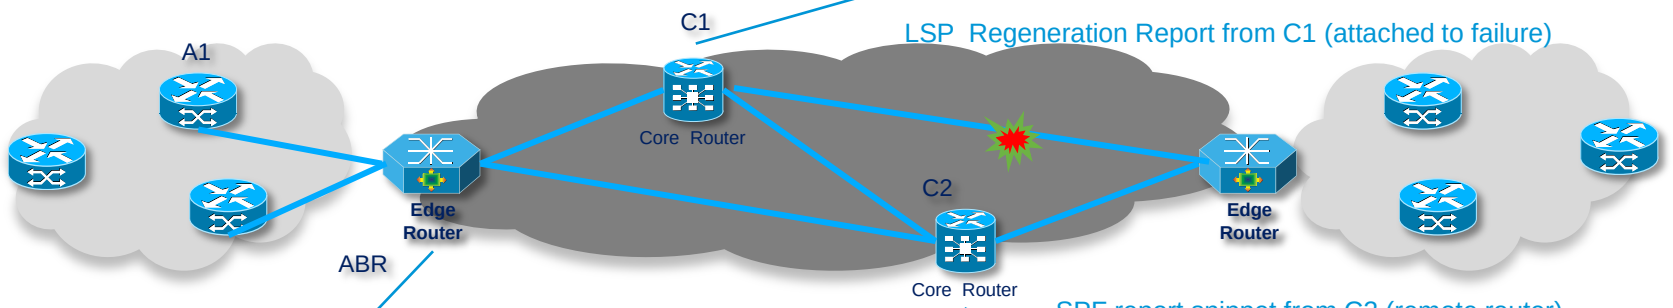

SPF report snippet from ABR – showing L2 SPF and trigger for its L1 LSP regen. into the attached L1 area

SPF report snippet from C2 (remote router)

```
Run: 80        Topology: 0              Level: L2               Type: Full
    Trigger: Feb 16 14:39:55.173           Trigger:  ll lp
                  Wait: 100                Start: 102              Duration: 1
    Trigger LSP:   0020.0203.2002.00-00  Seq: 50  Change-type: Modify  Time: Feb 16 14:39:55.173
<snip>
 LSP Processed:
       Id: 0020.0203.2002.00-00  Seq: 50  Change-type: Modify  Recv-Time: Feb 16 14:39:55.173
       Id: 0000.0000.0030.00-00  Seq: 2337  Change-type: Modify  Recv-Time: Feb 16 14:39:55.239
```

```
 Run: 111       Topology: 0             Level: L2               Type: Full
    Trigger: Feb 16 14:40:06.500           Trigger:  ll lp
                  Wait: 100                Start: 102              Duration: 1
    Trigger LSP:   0020.0203.2002.00-00  Seq: 50  Change-type: Modify  Time: Feb 16 14:40:06.500
<snip>
 LSP Processed:
       Id: 0020.0203.2002.00-00  Seq: 50  Change-type: Modify  Recv-Time: Feb 16 14:40:06.500
       Id: 0000.0000.0030.00-00  Seq: 2337  Change-type: Modify  Recv-Time: Feb 16 14:40:06.570
 LSP Regenerated:
       Id: 0001.0001.0001.00-00  Seq: 48  SPF: 111  Lvl: L1  Trigger:  ia   Regen-Time: Feb 16 14:40:06.657
```

A Case Study – SPF Event Summary

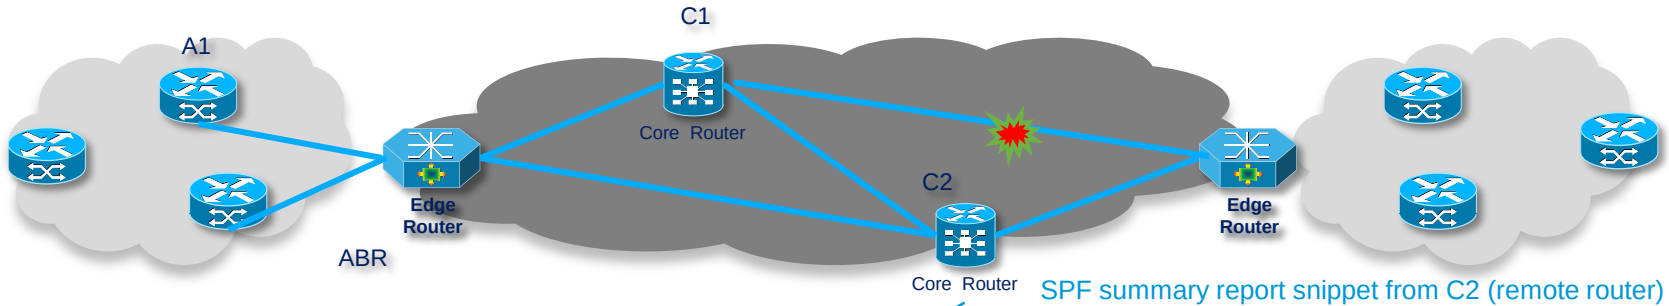

SPF summary report snippet from C2 (remote router)

```
Legend:
SPF                     - ISIS process level SPF counter.
Trigger Time            - Absolute time when the SPF was triggered (in mmm dd hh:mm:ss.msec).
Dur                     - Duration of the ISIS SPF computation (msecs).
Type                    - Type of SPF run.
Critical/High/Medium/Low - Priority based on the configure prefix prioritization policy.
For each priority       - Total prefixes affected and time taken for their programming (msecs) for IP and MPLS.
^  no route change #  threshold exceeded ~ incomplete data * collection pending

Reporting SPF Events for ISIS Instance : 1
```

SPF	Trigger Time	Dur	Type	LSPs	Critical	High	Medium	Low
71	Feb 16 14:04:14.864	9	PRCL	0	0 / - / -	3 / 116 / 125	3 / 117 / 126	64 / 122 / 127
73	Feb 16 14:05:43.768	1	FULL	2	3 / 107 / 116	3 / 107 / 118	3 / 107 / 119	6 / 108 / 121
74	Feb 16 14:24:30.108	0	PRCL	1	0 / - / -	0 / - / -	0 / - / -	1 / 107 / 112
75	Feb 16 14:24:34.978	1	FULL	2	3 / 107 / 118	3 / 107 / 120	3 / 108 / 122	5 / 108 / 125
76	Feb 16 14:28:50.800	1	FULL	2	3 / 107 / 116	3 / 107 / 118	3 / 107 / 119	6 / 108 / 121
77	Feb 16 14:37:36.491	0	PRCL	1	0 / - / -	0 / - / -	0 / - / -	1 / 106 / 112
^78	Feb 16 14:37:44.627	0	FULL	1	0 / - / -	0 / - / -	0 / - / -	0 / - / -
79	Feb 16 14:37:45.075	1	FULL	1	3 / 107 / 117	3 / 108 / 119	3 / 108 / 121	5 / 108 / 125
80	**Feb 16 14:39:55.173**	**1**	**FULL**	**2**	**3 / 107 / 117**	**3 / 107 / 118**	**3 / 108 / 119**	**6 / 108 / 121**

- Provides high level snapshot of SPF events – their impact on routes and convergence times
- Also identifies events where “threshold” has exceeded.

A Case Study – SPF Convergence Report

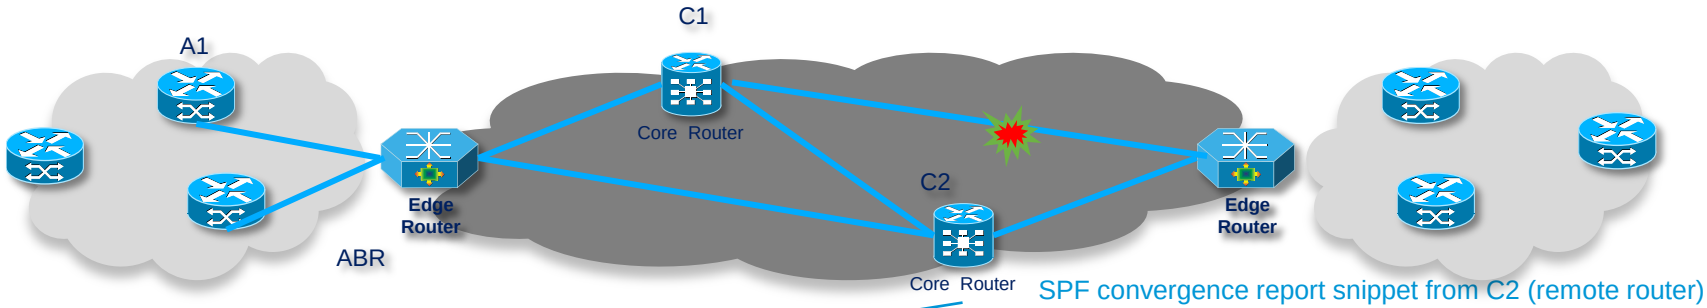

SPF convergence report snippet from C2 (remote router)

```
 Run: 80         Topology: 0              Level: L2               Type: Full
    Trigger: Feb 16 14:39:55.173             Trigger:  ll lp
                  Wait: 100                  Start: 102               Duration: 1
    Trigger LSP:    0020.0203.2002.00-00  Seq: 50  Change-type: Modify  Time: Feb 16 14:39:55.173
    Node Stats:     Added: 0                 Deleted: 0               Modified: 2
                    Reachable: 40            Unreachable: 0           Touched: 40
    Timeline Summary:
        Priority: Critical
            Route Count:                  Added: 0            Deleted: 0          Modified: 3
            IP Route Program Time:        Min: 107(0/0/CPU0)           Max: 107(0/0/CPU0)
            MPLS Label Program Time:      Min: 117(0/0/CPU0)           Max: 117(0/0/CPU0)
        Priority: High
            Route Count:                  Added: 0            Deleted: 0          Modified: 3
            IP Route Program Time:        Min: 107(0/0/CPU0)           Max: 107(0/0/CPU0)
            MPLS Label Program Time:      Min: 118(0/0/CPU0)           Max: 118(0/0/CPU0)
        Priority: Medium
            Route Count:                  Added: 0            Deleted: 0          Modified: 3
            IP Route Program Time:        Min: 108(0/0/CPU0)           Max: 108(0/0/CPU0)
            MPLS Label Program Time:      Min: 119(0/0/CPU0)           Max: 119(0/0/CPU0)
        Priority: Low
            Route Count:                  Added: 0            Deleted: 1          Modified: 5
            IP Route Program Time:        Min: 108(0/0/CPU0)           Max: 108(0/0/CPU0)
            MPLS Label Program Time:      Min: 121(0/0/CPU0)           Max: 121(0/0/CPU0)
<snip>
```

- Provides details on router-wide route update time, trigger details, statistics, fastest/slowest LCs, etc.

A Case Study – Leaf Network Events

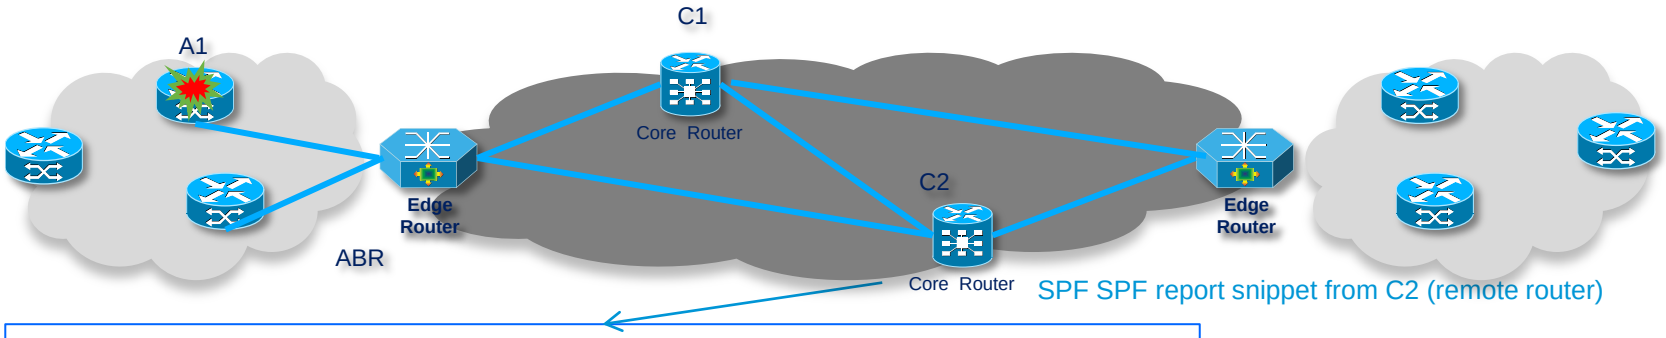


```
<snip>
Priority: Critical
       Route Count:                  Added: 0            Deleted: 10         Modified: 0
       IP Route Program Time:        Min: 107(0/2/CPU0)            Max: 108(0/0/CPU0)
       MPLS Label Program Time:      Min: 116(0/1/CPU0)            Max: 117(0/0/CPU0)
       Details:                     Start               End           Duration
             ISIS:                    102                102                 0
             RIBv4-Enter              102                102                 0
             RIBv4-Exit               106                106                 0
             RIBv4-Redist             106                107                 1
             LDP Enter                107                107                 0
             LDP Exit                 108                108                 0
             LSD Enter                108                108                 0
             LSD Exit                 112                112                 0
             LC Details(IP Path):
               S 0/0/CPU0             108                108                 0
                 0/1/CPU0             108                108                 0
               F 0/2/CPU0             107                107                 0
                 0/3/CPU0             108                108                 0
             LC Details(MPLS Path):
               S 0/0/CPU0             114                117                 3
               F 0/1/CPU0             114                116                 2
                 0/2/CPU0             113                116                 3
                 0/3/CPU0             114                117                 3

       Leaf Networks Added:
           -
       Leaf Networks Deleted:
           6.0.0.0/24                 6.0.1.0/24                 6.0.2.0/24
           6.0.3.0/24                 6.0.4.0/24                 6.0.5.0/24
           6.0.6.0/24                 6.0.7.0/24                 6.0.8.0/24
           6.0.9.0/24
```

- Provides (throttled) logging of delete/add of leaf networks
- Also internal convergence timeline for the router

Insights into Routing Convergence

- Characterization of convergence times for critical/high priority prefixes
  - How many %age converged in sub-second periods?
  - Impact on customer SLA
- Type 1-2 LSA (or LSP) flooding/propagation delays see in the network
- Data to monitor & analyze impact of network changes on convergence
  - How different routers reacted during a period of churn?
- Diagnostics mode automatically triggered and additional debug data collected when update times exceed specified threshold
- Detailed convergence data collected and archived for post-mortem analysis of critical and high impact failures
- Analysis on effectiveness of SPF & LSA/LSP timers
- Integration with intelligent offline tool that is topology aware to gather data for end-to-end network convergence

RCMD features available on IOS-XR Release 4.2.0 onwards on Cisco CRS1/3, ASR9000 and XR 12000 platforms

Conclusion - RCMD Overview

- Challenges exist TODAY in Monitoring and Analyzing Routing Convergence in production networks
- RCMD is a tool that reports data related to Routing Convergence
  - Provides "in-router" view of convergence events
  - Lightweight and always-on – non convergence impacting
  - Persistent – archived for use after hours/days
- Runs in Monitoring & Diagnostics modes
- RCMD reports provide detailed data points related to router convergence
- Leverage RCMD reports for network-wide convergence analysis
  - Impact to service SLAs
  - Impact of network design changes & growth

Thank you.