

MPLS Traffic Engineering with Auto-Bandwidth: Operational Experience and Lessons Learned

Mohan Nanduri
Mark Kasten
Naoki Kitajima

Microsoft
February 29, 2012

Agenda

- Offline vs. Auto-Bandwidth
- Auto-Bandwidth basics and overview
- Operational Experience
- Lessons Learned
- Configuration

MPLS LSP Bandwidth Adjustment

- Few ways to accomplish this:
 - Offline calculation
 - Calculating paths using offline tool
 - Periodically adjusting bandwidth value based on measured 95 percentile traffic plus adding additional buffer for growth
 - Auto-Bandwidth
 - Rely on devices to measure traffic and adjust bandwidth

MPLS Auto-Bandwidth Terminology

- Adjust-interval
 - Frequency at which LSP bandwidth is adjusted
- Adjust-threshold
 - Change in LSP bandwidth utilization percentage that triggers new LSP calculation
- Minimum and Maximum Bandwidth
 - Minimum and Maximum amount of bandwidth
- Overflow and Underflow
 - Number of consecutive bandwidth overflow/underflow samples before triggering bandwidth adjustment
- Monitor Bandwidth
 - Monitor usage of bandwidth without triggering adjustment

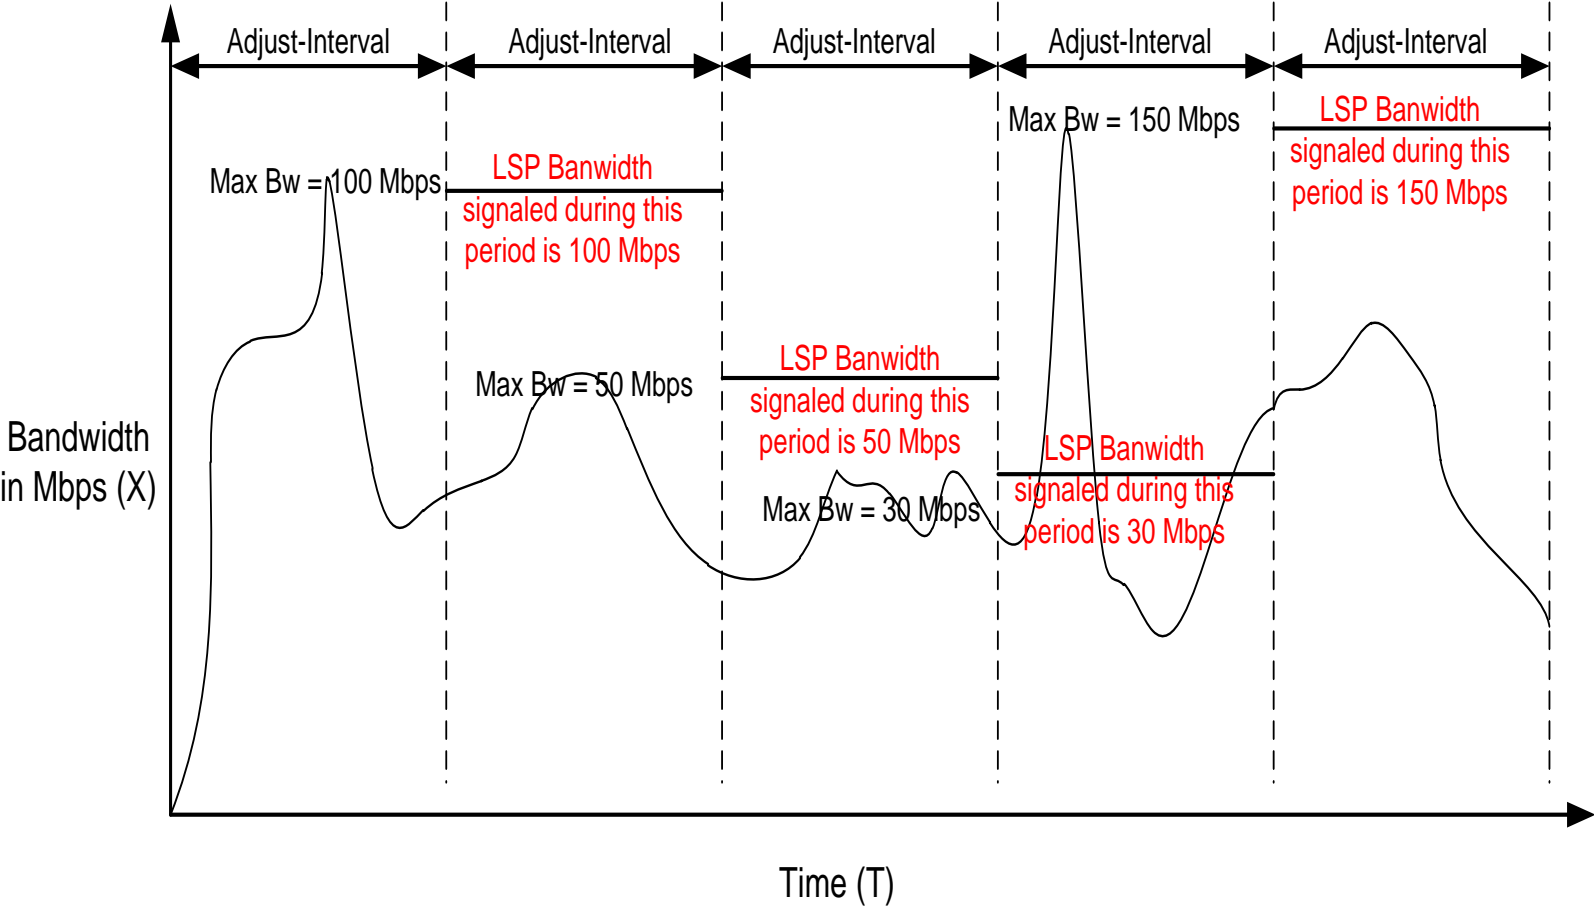
MPLS Auto-Bandwidth Overview

- Automates the process of monitoring and adjusting LSP bandwidth
- MPLS auto-bandwidth measures the traffic flows through the LSP, adjusting the bandwidth based on measured traffic and defined parameters on per-LSP basis
- LSPs are treated independently to adjust bandwidth without regard for any adjustments previously made or pending for other LSPs

MPLS Auto-Bandwidth Overview

- Auto-Bandwidth Calculation mechanism:
 - Bandwidth adjustment is based on per-LSP measured traffic utilization and chooses the highest value during the adjust interval
 - At the end of each interval the highest measured bandwidth value is compared with the existing LSP bandwidth value
 - If the delta between existing LSP bandwidth and highest measured bandwidth value is equivalent or greater than adjust-threshold value, LSP tries to re-signal with the new bandwidth, otherwise the LSP bandwidth remains the same
 - Make-before-break mechanism kicks into place and LSP re-signals for the new bandwidth. Once it's successful in setting up the LSP, the maximum bandwidth value is reset to zero
 - If the new bandwidth cannot be setup, LSP uses the existing bandwidth value
 - The above process is repeated for the next adjust-interval


MPLS Auto-Bandwidth With Adjust-Interval



MPLS Auto-Bandwidth With Adjust-Interval

- The highest measured bandwidth value is 100 Mbps for the given adjust-interval period, LSP is signaled with 100 Mbps and its reset to zero
- As the new adjust-interval starts, the highest measured bandwidth value is 50 Mbps. LSP is signaled with 50 Mbps if the adjust-threshold value is met

MPLS Auto-Bandwidth With Overflow



MPLS Auto-Bandwidth With Overflow

- If the highest measured bandwidth value is higher than the current LSP signaled bandwidth value
- The delta between highest measured bandwidth value and current LSP signaled bandwidth value is greater than or equal to configured adjust-threshold value
- The number of occurrences must reach the configured overflow count consecutively. Overflow counter resets to the configured overflow value if the bandwidth decreases during one the sampling cycles
- LSP is signaled with the new bandwidth value before the configured adjust-interval
- Adjust-interval timer is reset to configured value.

Benefits of Auto-Bandwidth

- Network can react faster to sudden burst of traffic in near real-time and not rely on manually intervention
- Effective use of bandwidth resources by minimizing the over-subscription/padding of LSP bandwidth
- Maximizes the usage of available bandwidth and optimizes the network effectively to use more preferred, shorter latency, paths first.

Why auto-bandwidth for us?

- Using offline bandwidth calculation based on historical traffic data was slow, cumbersome, and didn't support real time growth
- Varying services with different traffic profiles
- Lagging behind bandwidth adjustments
- Not reacting to sudden spikes in traffic shifts

Operational Experience

- Some skepticism initially letting the device make decisions and automatically adjust bandwidth
- Auto-bandwidth deployed since early 2009
- Phased rollout of auto-bandwidth globally
- Extensive analysis was performed on network after auto-bandwidth was introduced to optimize
 - Significant churn by smaller sized LSPs
 - Tiered LSP model was introduced to reduce churn
- RSVP interface subscriptions were modified to account spike in traffic
- Sudden spike in traffic is still an issue
 - Automatic LSP creation and deletion using [event scripts](#) introduced
- Overall Worked fairly well so far

Lessons Learned

- Significant amount of time was spent to understand internals, was tough to get information initially from vendor
- Many hours were spent testing and validating, scripts to validate saved lot of time
- Testing all the available knobs thoroughly helped, even though we did not plan to use them right away. We ran into multiple issues with these knobs
- Thorough planning and training made the deployment very smooth
- Post deployment correlation scripts to check traffic stats and MPLS bandwidth helped to check any anomalies

Lessons Learned

- Show commands lacks historical bandwidth adjustment information, have to glean information from multiple log files
 - Multiple enhancements have been submitted to display information to make troubleshooting easier, to name a few:
 - Show previous bandwidth adjusted values
 - Minimum/maximum bandwidth ever used
 - Counter to keep track of number of bandwidth adjustments
 - Ran into some issues early on
 - Smear timers not spacing LSPs properly
 - Bloated traffic statistics used to setup LSPs
- Based on configured values, smaller LSPs were causing churn and analyzing a weeks worth of data showed more than 5000 events

Configuration

```
groups {
  autobw_lsp_setup {
    protocols {
      mpls {
        label-switched-path <*> {
          ldp-tunneling;
          soft-preemption;
          admin-group include-any [ gold];
          adaptive;
          auto-bandwidth {
            adjust-interval 1200;
            adjust-threshold 10;
            minimum-bandwidth 1m;
            maximum-bandwidth 2g;
          }
        }
      }
    }
  }
  autobw_lsp_setup_1M_profile {
    protocols {
      mpls {
        label-switched-path <*> {
          ldp-tunneling;
          admin-group include-any [ gold ];
          adaptive;
          auto-bandwidth {
            adjust-interval 1200;
            adjust-threshold 50;
            minimum-bandwidth 1m;
            maximum-bandwidth 50m;
          }
        }
      }
    }
  }
}

label-switched-path rtr1-to-rtr2 {
  apply-groups autobw_lsp_setup;
  to 10.1.1.2;
  admin-group exclude [bad];
}

label-switched-path rtr1-to-rtr3 {
  apply-groups autobw_lsp_setup_1M;
  to 10.1.1.3;
  admin-group exclude [bad];
}

mpls {
  statistics {
    file mpls_stats.log size 1m files 10;
    interval 60;
    auto-bandwidth;
  }
}
```


Thank You