CASE STUDY

# Keeping Your Fabric Healthy and Performing Optimally with Mellanox UFM®

## Background

A recent IBM supercomputing project required an environment that included over 1000 nodes with both GPFS GSS servers, other GPFS storage servers, a main production cluster, and a small test cluster. In order to ensure the healthiest and highest performing interconnect for the project, IBM and the customer chose to deploy Mellanox Unified Fabric Manager (UFM)® for data center management.

## The Solution

UFM is a powerful platform for managing scale-out Ethernet and InfiniBand computing environments. UFM enables data center operators to efficiently provision, monitor, and operate the modern data center fabric. This enables more efficient troubleshooting and higher utilization of fabric resources.

IBM achieved two primary objectives by installing UFM:

- Obtained a real-time source for monitoring fabric performance and traffic characteristics, leading to optimization of the fabric performance
- Quickly identified and corrected routing issues, improving overall fabric performance and availability

## Monitor and Optimize Performance

First, the site used UFM's logical model to implement service-based modeling. This grouped the data center's physical servers into logical entities, from each of which monitored data was then aggregated to provide a better understanding of the behavior of the respective group (such as GPFS GSS servers, other GPFS storage servers, its test cluster, and its production cluster).

UFM's advanced monitoring engine aggregated the performance information and correlated it with the logical entities. This produced specific data that enabled IBM to analyze health and performance issues more granularly since they were associated with particular services.

## Detect and Reduce Congestion

Congestion is often difficult to identify, as it is not a physical error and does not issue typical failure alerts. UFM has unique capabilities that identify and notify the customer of congestion in the fabric.

IBM employed UFM's congestion map to receive alerts when congestion was present and to isolate its location. Almost immediately, IBM identified congestion on the congestion map, indicating a possible routing problem in the fabric.

Using UFM's dashboard and detailed tables, IBM quickly determined the nodes most affected by congestion, enabling them to analyze the problem and identify the potential for a routing issue.

Upon further analysis of the fabric routing and traffic patterns, UFM was able to find imbalances in the routing scenarios and to identify unused or idle ports.

The analysis yielded a set of correction actions that were implemented in the subnet manager routing configuration. By analyzing the fabric again with UFM, it was readily apparent immediately after corrective action was taken that while congestion was not completely eliminated, the main issue standing in the way of the cluster's performance had been corrected. One could see a more acceptable congestion pattern,


## OVERVIEW

*Mellanox UFM provided the right tools to solve the problems, allowing IBM to provide superior application performance, significantly reduce issue-handling time, and achieve the desired level of fabric availability for the running applications.*

less time spent in congestion, and better use of the available ports.

## Monitor and Debug Techniques

Many UFM tools were utilized in the process, including detailed congestion views, port level performance analysis, logical representation of server groups, and health and performance graphical monitoring sessions. To maintain brevity, the main result views of the real fabric are presented below:

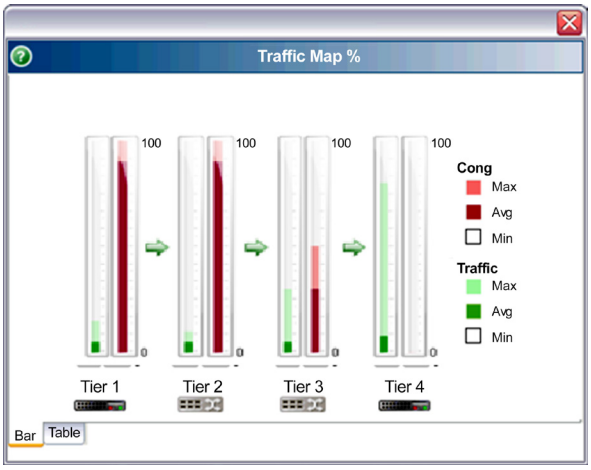


*Before:* Heavy congestion on Tiers 1 to 3 indicating potentially non-optimal routing. The average congestion is nearly the same as the maximum congestion and the number is large, which indicates that the network is running at maximum congestion most of the time. Traffic is low and apparently applications cannot send data across the fabric.

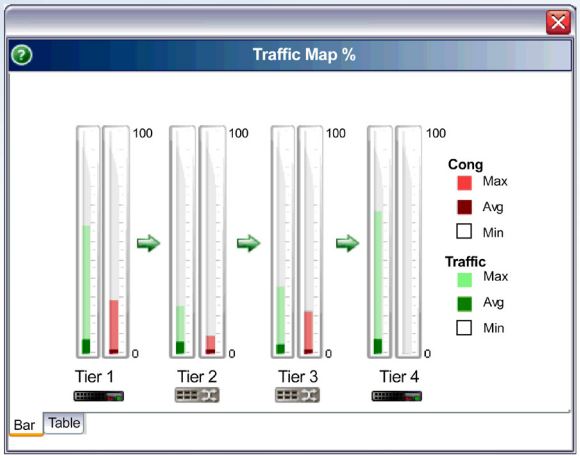


*After:* After corrective action was taken, the congestion pattern was more as expected, and the maximum congestion is a much lower percentage of the average. Maximum values of sent traffic are much higher, indicating that the servers that are trying to communicate are able to do so freely.

## Conclusion

IBM ultimately used UFM to successfully analyze performance trends per service group, and to quickly identify performance issues such as congestion. Mellanox UFM provided the right tools to solve the problems, and then supplied immediate and ongoing monitoring of the fabric condition.

This allowed IBM to provide superior application performance, significantly reduce issue-handling time, and achieve the desired level of fabric availability for the running applications.

Mellanox TECHNOLOGIES

350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com



15-1372B Rev 1.0