



Dell Active Fabric Manager 1.5 Features

Whitepaper

March 2013

© 2013 Dell Inc. All Rights Reserved. Dell, the Dell logo, and other Dell names and marks are trademarks of Dell Inc. in the US and worldwide.

This document is for informational purposes only and may contain typographical errors. The content is provided as is, without express or implied warranties of any kind.

Contents

1	Introduction	4
2	Active Fabric.....	5
3	Active Fabric Manager.....	7
3.1	OVF Compliant Virtual Machine.....	7
3.2	AFM Application	7
3.2.1	Design phase	7
3.2.2	Build phase.....	8
3.2.3	Provisioning phase	8
3.3	Deployment and Validation of Interlink fabric	8
3.3.1	Bare Metal Provisioning (BMP).....	8
3.3.2	Validation Status	9
3.4	Runtime phase	9
3.5	Open flexible architectures	10
3.6	New features in AFM 1.5.....	10
3.7	Virtual Link Trunking (VLT).....	10
3.7.1	Significant advantages of VLT.....	10
3.8	AFM for VLT	11
3.8.1	VLT modes	11
3.9	Design Templates.....	12
3.10	Audit Trail.....	13
3.11	Versatility in Operation.....	13
3.12	Unique benefits of the AFM.....	14
3.13	Summary.....	14

1 Introduction

This whitepaper discusses the significance of an automation application for the deployment of Dell's Active Fabric solutions. Specifically, this paper highlights the unique features of Dell Active Fabric Manager (AFM) 1.5 that automate the design and deployment of fabric for the next generation data centers.

Managing the data center for high availability is becoming significant and essential. The management application has to take care of creating an optimal design and enabling automated deployment, alert the user on active alarms, analyze performance and assist in troubleshooting issues. As the next generation data centers emerge with the latest concepts and evolving standards, optimizing the performance of compute, storage and networks, managing the infrastructure have become a challenge. With virtual machines housing business critical applications as well as transition from 1G to 10GE and beyond, network management becomes complex and crucial to meet the service level agreement. Extending the high availability and resiliency of efficient flat layer-2 network with auto provisioning and management for Virtual link Trunk deployed network is one of the new features with AFM 1.5.

2 Active Fabric

Active Fabric is the umbrella term that refers to the suite of fabric offerings from Dell. Over the last 2+ years, we have introduced some very unique cost-effective, high performance and massively scalable fabrics catering to the needs of large datacenters. As more and more applications are getting virtualized and there is a rapid acceleration towards deploying and consuming services on the cloud, the need for fabrics is becoming more and more relevant for datacenters of all sizes.

At Dell, we are enabling our customers by offering the entire portfolio of infrastructure elements including server, storage and networking elements. Our Active fabric is targeted at offering our customers to easily scale racks and racks of servers and storage equipped with 1 Gb, 10 Gb and 40 Gb interfaces. We have made significant strides since we introduced the distributed core fabric. There are a number of switching elements in our portfolio that are part of the Active Fabric offerings which will continue to evolve with the release of newer switches.

Generally, the fabric woven around these nodes built on the CLOS architecture, orchestrates the seamless live migration of VMs with consistent flow of traffic across these servers in east-west and north-south patterns. These next generation networks have to be optimally designed and effectively managed to meet the end-to-end performance requirements.

As the number of nodes in a fabric keep increasing, effective management becomes quite complex for the Network/System admins. Node image upgrade management, Configuration management, Performance management, Fault management forms the significant complexity in managing and monitoring these fabric networks.

Modularity is the key and one of the main benefits in Dell's Active Fabric. Resiliency in operation is ensured by means of redundant nodes with its various interlink passing traffic through the fabric. Any link or node failure has no major impact for the business critical traffic.

The network fabric is deployed with fixed form factor switches promoting fast, fat and flat network. With the Dell networking nodes, operating power requirement is significantly reduced in the Spine-Leaf architecture, the overall AC/DC power required is very less improving greener IT and HVAC requirement gets minimal.

Dell Networking while delivering the latest trends and pioneering the Ethernet fabrics has laid out a clear and robust design for next gen data centers. With the clear understanding of unique customer requirements, customizable and optimized network design for the end to end solution, Dell has rolled out AFM application to design, build and manage the fabric.

Figure 1 shows the typical Spine-Leaf interlinks connecting all the leaf nodes from all the Spine nodes.

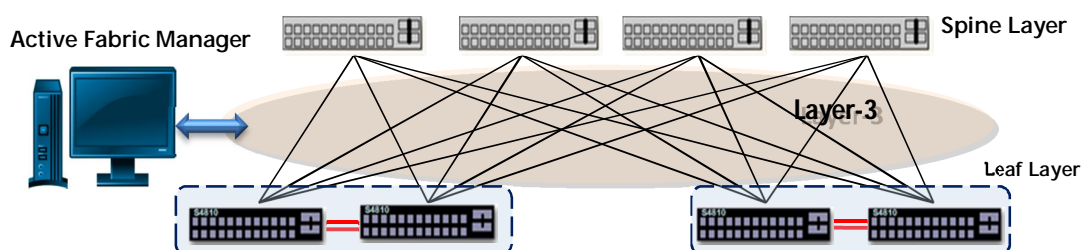


Figure 1. Typical Spine-Leaf CLOS Network

Based on the customer's unique requirements, AFM designs the number of downlink access ports (Server-connect ports) which connect from the Leaf or Top-of-Rack node, with specific links interconnecting to the Spine nodes. The inter-link handles the bandwidth for East-west and North-South traffic. This could be of varying ratio of 1:1 non-blocking line-rate capacity of downlink to Interlink traffic, and be flexible to 6:1 also. The oversubscription is generally based on the applications and the services deployed on the servers connected to the underlying ToR Switches.

3 Active Fabric Manager

The Dell Active Fabric Manager (AFM) application is built on Core Java (J2SE) framework with Spring MVC (Module View Controller) for modularity in implementation. It supports multiple modules and is a generic web based Client-Server model; with integrated modules serving North and South bound interfaces with middleware services and Application layer.

3.1 OVF Compliant Virtual Machine

AFM can be deployed as a virtual appliance with the Open Virtualization Format (OVF) template on a VMware ESX virtual machine or workstation. After completing the OVF wizard and deploying the template, you can start AFM VM by using "Power On" and opening the "Console" to log in with default credentials. Access the VM through the standard browser to perform the initial configuration. You must accept the license agreement before proceeding to the System and Service Protocols to configure the Syslog, TFTP, FTP and DHCP Servers.

3.2 AFM Application

The base modules in the AFM application are:

1. Design Module
2. Build Module
3. Provisioning Module
4. Performance Management Module
5. Fault Module
6. Node manager view
7. Backup/restore, FTOS image upgrade, Audit trail

3.2.1 Design phase

Designing the data center fabric is made easy with the complex design process handled by the AFM application. You must specify the number of server ports and the uplink ports for the upstream traffic. The design wizard also provides the option of selecting the ports for future requirement to meet the scalability needs of pay as you grow model. AFM assists the user to setup layer-3 fabric or layer-2 fabric (with VLT). The initial screen gives the option of selecting distributed fabric or VLT deployment. The flexible types of deployment can be chosen to have Z9000 and S4810 Dell Networking Switch as Spine and leaf nodes forming the fabric, also supporting PowerEdge M1000e Chassis with MXL Blades. Line rate or oversubscribed traffic demands the stature of multi-tiers in tandem to connect the underlying business-critical servers to the external world.

The number of uplink, downlink ports, bandwidth, over subscription rate along with the type of fabric are the major input parameters in design phase. The output of this module is Network topology and the detailed wiring plan. Similarly for VLT, the number of ports and the type of fabric decides the topology. MXL switches in the Dell M1000e Blade Chassis function as ToR Switch in the VLT deployment.

AFM automates the task of Network design, Optimal validation, Configuration deployment , Performance and Fault management for the Active Fabric.

3.2.2 Build phase

The network design output is generated by the design wizard with Network topology map, tabular and graphical wiring plan with Interlink connectivity chart and the port map chart of each node. The detailed BOM and the charts can be exported and printed for ease of use. The distributed core or VLT configuration for all the nodes along with the start-up configuration is generated in the build phase.

3.2.3 Provisioning phase

Pre-deployment, deployment and validation of the network design forms part of the provisioning phase. MAC Address, IP address schema, Interlink OSPF protocol, uplink protocol parameters and user-defined VLAN's are furnished through the pre-deployment wizard before deployment. Each of the process status is indicated as the configuration is deployed on the nodes and finally the network design is validated.

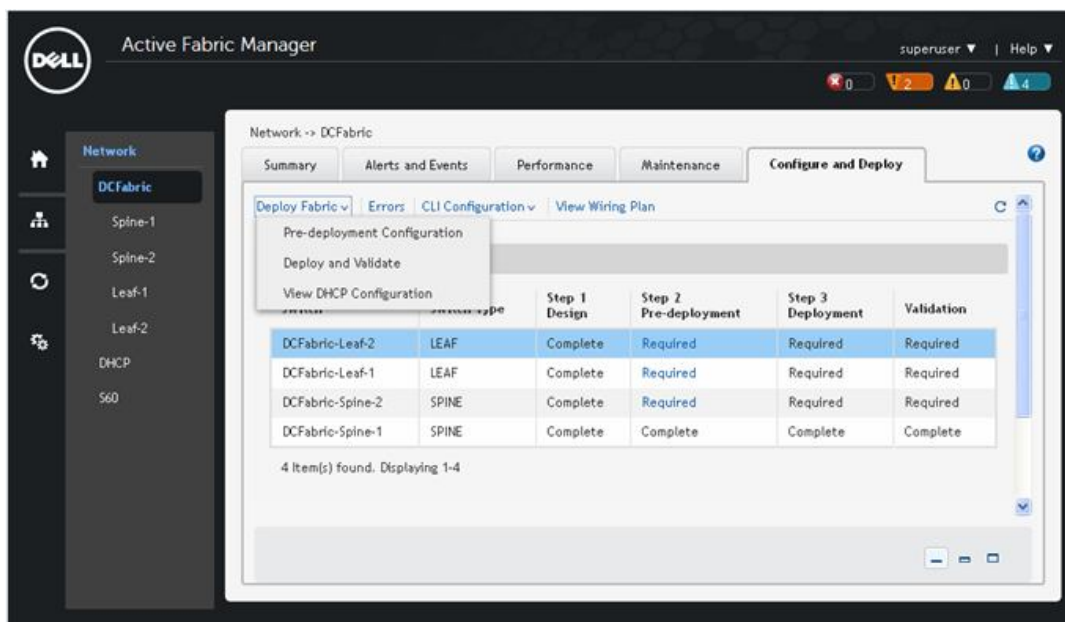


Figure 2. Active Fabric Manager Screen

3.3 Deployment and Validation of Interlink fabric

After the Fabric design validation, the cabling of the QSFP/SFP+ ports with Twinax Direct attach cable (DAC) or optional Optics on OM3 fiber can be decided based on the detail summary.

3.3.1 Bare Metal Provisioning (BMP)

Open Automation framework, delivering whole new level of architecture for the Dell Networking Switches has the BMP Jumpstart feature enabling the node automation to upgrade the FTOS image from a TFTP Server after getting the assigned IP address from the DHCP Server.

Firmware upgrading and the Configuration deployment for multiple nodes in an Active Fabric would be otherwise a cumbersome process taking more time and less efficient.

With AFM deploying all these nodes to get the Firmware image (FTOS) and the appropriate configuration through the BMP process, the overall operation is automated and implemented effectively improving the convergence time for fabric bring up.

3.3.2 Validation Status

Validation of all or selective nodes of the fabric with configuration deployment would be indicated with the relevant status details. Any error in the deployment and validation is displayed with response action.

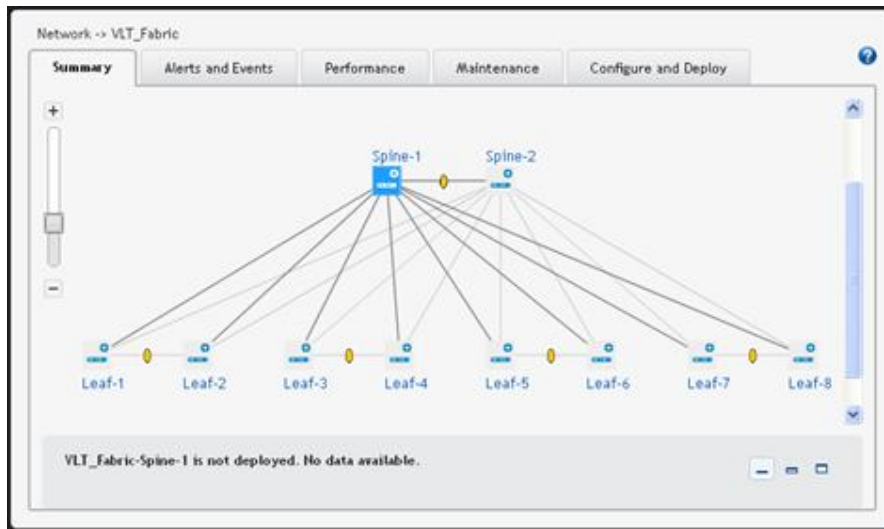


Figure 2. Active Fabric Manager Screen

3.4 Runtime phase

Fault management, Performance management, Node manager View, Audit trail, Periodic validation and Maintenance modules form part of the Runtime phase.

Health status with unique alarm is indicated in dashboard. Alerts and events with current and historical data for each network and fabric are displayed with details. Filter options with various parameters could search the specific alarm or event within its operational period.

Top 10 usage of Bandwidth, CPU and memory for a device or a fabric are displayed in the dashboard. With real time polling, the active utilization can be viewed in tabular/Line chart/graphical view. The granular traffic pattern could be zoomed with different time intervals for user requirement. Also customized performance reports could be created for a particular metric for the historical data. With the node view, the user can select and view the individual ports for the throughput and error rate.

AFM does periodic validation of deployed fabric and highlight to the user, if there is any mismatch in node, wiring and configuration change.

Scheduled FTOS upgrade for the complete or partial fabric could be done through the maintenance wizard. With the maintenance scheduler, the user could avoid complete downtime of the fabric.

3.5 Open flexible architectures

Configuration of all the participating nodes in the Spine-Leaf Active fabric network is ensured after validating the Interlink connectivity requirements. With OSPF enabled on the Interlinks, the Server Subnets are advertised to all other ToR Switches for the inter-vlan reachability and optimally tweaked for fast convergence for the predominant east-west traffic in today's data centers. For the layer-2 fabric, the user defined VLAN's with the associated VLT provisioning is totally automated.

Uplinks to other Data center Clusters, or to the upstream ISP's with interoperable protocols OSPF, iBGP or eBGP can be configured to have seamless flow of North-South traffic across the Data Centers.

3.6 New features in AFM 1.5

1. OVF support for AFM deployment
2. Automation of VLT in Z9000 and S4810 Switches
3. MXL in M1000e Chassis connected to VLT
4. Design templates
5. Audit trail

3.7 Virtual Link Trunking (VLT)

VLT is a feature implemented in FTOS, ensembles two physical switches to represent as a single logical switch. With physical links as a port-channel connecting two individual switches configured with VLT would logically group and resemble it as single entity only for the access switches which connects to the VLT domain. Both the VLT peers have their own configuration, control and data planes. VLT permits the port-channels connected across two Switches and eliminates the need for Spanning tree protocol (STP) blocking state. With VLT, all the links connected to the access/ToR Switches are effectively utilized. The access device could be a Switch, Switch Stack, Server or any other device supporting LACP port-channels. High-availability and redundancy in the network is achieved by its duality in physical connectivity.

3.7.1 Significant advantages of VLT

1. Loop free connectivity in layer-2 domain
2. Faster Network convergence
3. High-availability and redundancy
4. Effective utilization of all the links
5. Link level resiliency
6. Active-Active Load sharing with VRRP.
7. Agility in VM Migration under VLT domain.
8. Unified access for virtualization, Web applications and Cloud computing
9. High performance for Big Data networks
10. Easier manageability of Network.

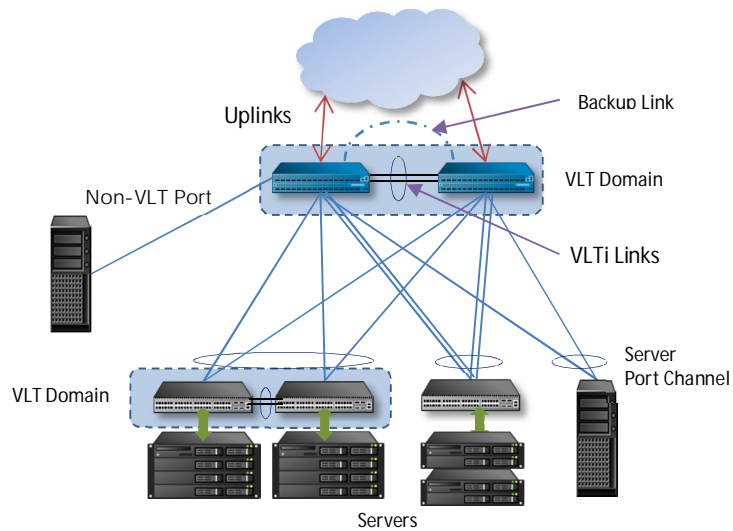


Figure 3. Virtual Link Trunking

3.8 AFM for VLT

Active Fabric Manager 1.5 has the VLT management capabilities incorporated for auto configuring the VLT peers and its ToR switches based on the user's requirement, besides managing the distributed core fabric. With auto configuring VLT domains, VLT Interconnect (VLTi) ports which connects the VLT peers within a VLT domain and the VLT port-channels, it is much easier for network administrators to manage the VLT domain through AFM.

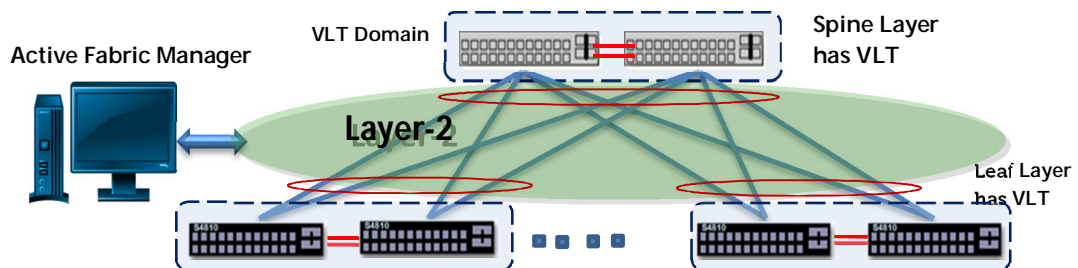


Figure 4. Multi VLT domains

3.8.1 VLT modes

AFM supports simple VLT in Spine model and mVLT in Spine and leaf model. With VLT in Spine only model, the leaf nodes could be of high performance S-series Switches or MXL Blade Switch integrated in PowerEdge M1000e Server Chassis.

mVLT deployed in both Spine and leaf nodes have Z9000 pair as Spine nodes and multiple S4810 pairs as Leaf nodes for large data centers. mVLT is typically deployed to extend the user defined VLAN's across multiple ToR nodes. Large data centers require same VLAN for migration of Virtual Servers spanning the physical servers connected to the ToR Switches, utilizing the layer-2 fabric infrastructure.

The following table illustrates the VLT deployment scenarios:

	Type-1 (1Gb)	Type-2 (10Gb)	Type-3 (10Gb)
Spine Switch	S4810 (2 Nos)	S4810 (2 Nos)	Z9000 (2 Nos)
ToR Switch	S55 or S60 (48 Nos)	MXL Blade 2 x 12 Chassis	S4810 (26 Nos)
O/S Ratio	2.4 : 1	3 : 1	6 : 1
Uplink Ports	4 x 40 Gb	8 x 40 Gb	4 x 40 Gb
Server Ports	2304 (1 Gb ports)	240 (10 Gb ports)	1248 (10 Gb ports)

AFM besides designing the layer-2 network with VLT, manages the VLT Spine and leaf nodes for performance and fault management.

3.9 Design Templates

Another new feature in AFM 1.5 assists the user in creating multiple configuration templates which could be applied for any or group of the fabric nodes requiring a specific repeated configuration.

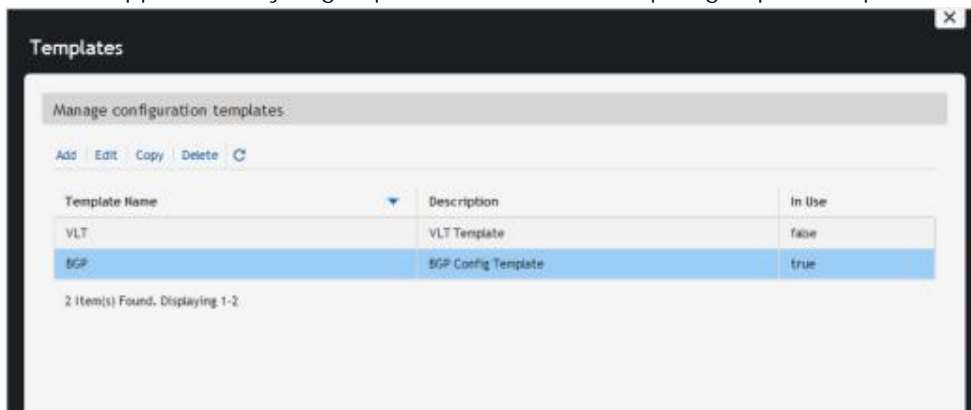


Figure 4 (a). Managing the Template

This process is made easy by associating the created templates to all or specific nodes to accelerate the adoption of Active fabric.

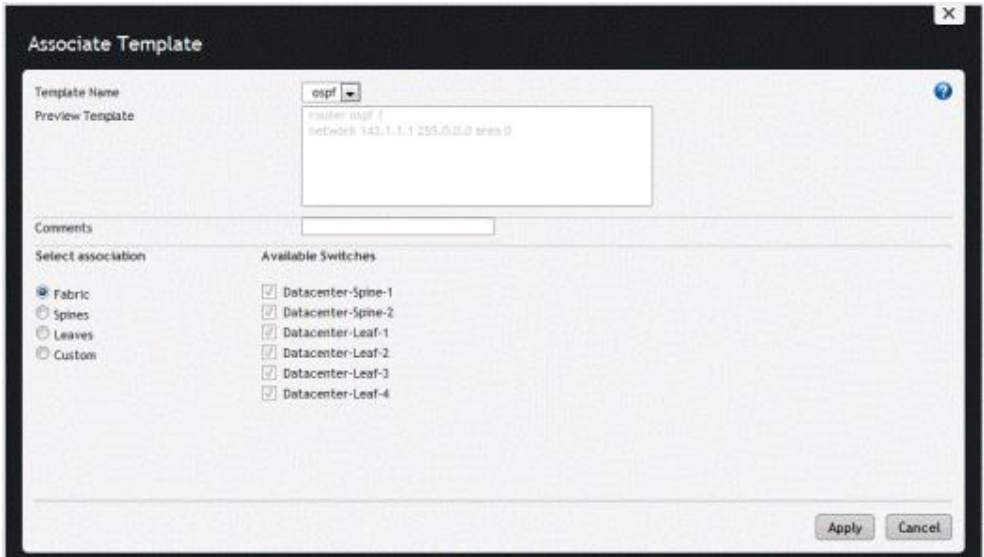


Figure 4 (b). Associating the Template

3.10 Audit Trail

The audit log with the login user name, date and time with the corresponding operation status is indicated in detail as shown in figure 5.

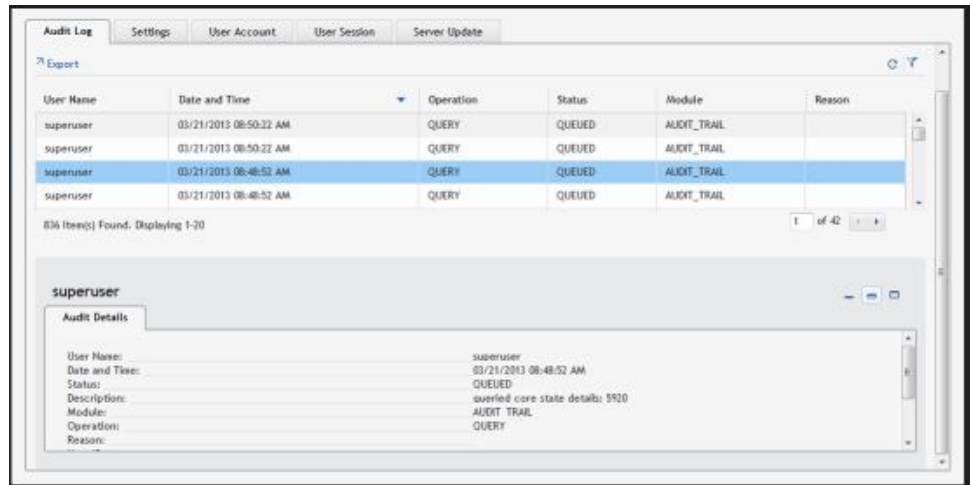


Figure 5. Audit Details

3.11 Versatility in Operation

Performance reporting for the fabric link utilization is in real time supporting historical data. Alarms and Events display the health of the fabric for the user to take proactive action. Any link failure gets critical alarm to be acknowledged by the user. The elegant user friendly dashboard sets the quick review of the alarms and the performance management statistics. Scheduler does the timely node backup, Performance Management data collection and Node software upgrade.

3.12 Unique benefits of the AFM

Feature	Benefit
Design templates for Layer-2 and Layer-3 fabric topologies	<ul style="list-style-type: none">• Accelerate the deployment of Distributed Core and VLT• Provide a single tool to design, document, deploy and manage complex network topologies
Automated Provisioning, Validation and Configuration	<ul style="list-style-type: none">• Task and time compression associated with deployment of complex fabrics• Near instant validation and rapid troubleshooting
Visualization & Usability Enhancements	<ul style="list-style-type: none">• Provide a rich end user experience• Establish a new way for customers to design and deploy complex networks
Package and distribute AFM as a holistic virtual appliance	<ul style="list-style-type: none">• Reduce the number of pre-requisites to install and use AFM• Easier end user experience
High Availability for AFM Server	<ul style="list-style-type: none">• No manual intervention to recover the AFM server in case of a failure
Northbound APIs for Ecosystem integration	<ul style="list-style-type: none">• Network abstraction for end users of higher level orchestration engines• Leverage AFM's functionality beyond the networking domain
Expanded Platform Support	<ul style="list-style-type: none">• Expanded support matrix – Z9000,S4810, S4820T,MXL,S60 and S55• Broader choice for customers
User Profiles	<ul style="list-style-type: none">• User profiles maps the individual user based role for relevant operation

3.13 Summary

Active Fabric Manager (AFM) is the ideal management application designed to manage the complex high density, low latency Active Fabric. Designing the network, validation, configuration, performance and fault management are the significant factors in managing the Next generation networks, and obviously Dell AFM does it with ease and with the power to do more.