

# Tellabs® ServiceAssured™ Upgrade

## With the Tellabs® 8800 Multi-service Router Series

Paymon Mogharabi  
 Manager, Product Line Management  
 Advanced Data Products

Today's telecom service providers face increasing pressure to meet conditions set forth in Service Level Agreements (SLAs) for system availability.

### Abstract

Multi-service routers deployed at the edge of the network are required to maintain continuous system availability while supporting more incoming customer traffic. The Tellabs® 8800 Multi-service Router Series is the industry's first carrier-class multi-service platform that meets the challenge of providing continuous system availability. Tellabs offers service providers a competitive advantage with Tellabs® ServiceAssured™ upgrades. With a Tellabs ServiceAssured upgrade, service providers can ensure that software-related maintenance upgrades do not adversely affect their existing revenue stream. With the Tellabs ServiceAssured upgrade, service providers can make changes to their network infrastructure with minimal disruption to their end-customer traffic. This paper defines the parameters that influence the network availability and looks at how Tellabs ServiceAssured upgrades can help address outage concerns during software maintenance. It discusses the objectives and requirements for performing a Tellabs ServiceAssured software upgrade and the steps taken to complete the software upgrade process. Finally, to truly appreciate the benefits of Tellabs ServiceAssured software upgrade, a comparison is made with the existing software upgrade solutions offered today by the vendors in the routing industry.

### Introduction

Service providers are increasingly offering specific availability and performance guarantees to their end-customers through SLAs. An SLA will protect end users against outages and performance degradations that are explicitly covered in

the agreement. Most SLAs cover areas such as availability, packet loss and latency. While packet loss and latency are important parameters within SLAs, the predominant condition set in SLAs today is system availability. Availability is the portion of total operating time that a network resource can be accessed successfully. This is the probability that the system is operating properly when it is requested for use. Service providers place a heavy emphasis on avoiding system unavailability or potential network downtime when designing their network infrastructure.

Network downtime can cost millions of dollars per hour in lost revenue to customers deploying mission-critical applications. Customers expect zero downtime and absolute availability from service providers. From a customer perspective, network downtime prevents them from focusing on their core business goals. Based on an Infonetics 2005 survey, the most important factor used by end customers to select service providers is availability. The amount of downtime as a result of an outage or system degradation has direct impact on the revenue stream for a typical Enterprise. Table 1 below provides a comparison of average hour and cost associated with the five vertical market sectors. The financial and manufacturing sector had the highest percentage of revenue lost due to network outage.

The impact of network downtime is particularly visible with edge routers given their strategic location within a network. An edge router is the aggregation point for incoming customer

|                            | Average | Financial | Health | Log    | Manufacturing | Retail |
|----------------------------|---------|-----------|--------|--------|---------------|--------|
| Total Hrs                  | 1,180   | 393       | 298    | 766    | 518           | 478    |
| Total Cost                 | \$222M  | \$42M     | \$32M  | \$154M | \$41M         | \$77M  |
| % of Revenue               | 16%     | 4%        | 2%     | 9%     | 5%            | 6%     |
| % of Cost From Outage      | 61%     | 59%       | 69%    | 69%    | 77%           | 65%    |
| % of Cost From Degradation | 39%     | 41%       | 31%    | 31%    | 23%           | 35%    |

Table 1. Comparison of Average Hours and Cost

traffic. Most edge routers strive to achieve higher densities to meet the growing end-customer traffic demands. In contrast to core routers, network architecture topologies do not factor in the same level of redundancy for their edge platforms as they do for their core routers. Many of today's core routers provide high availability through a dual connectivity model. While this is sufficient for the network core, it is not always economically or operationally feasible for edge router deployment. An edge router is more likely to be a single point of failure within a network. Compared to a core router, edge router downtime will have a much more significant impact on customer traffic. It is important for an edge router to maintain continuous system availability while forwarding customer traffic to the core. This is critical to a network operation center looking to maintain its network infrastructure while avoiding costly network downtimes.

System availability is based on redundancy and maintenance time. Redundancy is addressed through the use of redundant components or through higher MTBF numbers. Repair time or maintenance for the purpose of fixing defects or preventing future operational downtimes may require software upgrades. In theory, the number of times an upgrade is performed should not have any relevance to the availability time of the system. However, the majority of router products in the industry today cannot provide continuous system uptime while performing software upgrades. Tellabs ServiceAssured upgrades address the maintenance or repair time as it relates to the overall availability of the system.

In telecom, availability is measured in terms of “number of nines” or defects per million. This is equivalent to the percentage of time in a year that a system is functional and available for access by the user. With some service providers, the availability of a system is measured on the number of defects found within a large sample size. A Defect Per Million (DPM) value of 10 corresponds to a 99.999% reliability, equivalent to five minutes of downtime per year. While DPM is a useful measurement in scenarios where there is a partial network failure, the remaining portion of this paper will use the “number of nines” to measure system availability. In Table 2 below, a system with four nines is equivalent to a downtime of less than 53 minutes in one year.

| Availability | DPM (Defects Per Million) | Downtime Per Year (365x7x24) |
|--------------|---------------------------|------------------------------|
| 99.0         | 10,000                    | 3 Days, 15 Hours, 36 Minutes |
| 99.5         | 5,000                     | 1 Day, 19 Hours, 48 Minutes  |
| 99.9         | 1,000                     | 8 Hours, 46 Minutes          |
| 99.95        | 500                       | 4 Hours, 23 Minutes          |
| 99.990       | 100                       | 53 Minutes                   |
| 99.999       | 10                        | 5 Minutes                    |
| 99.9999      | 1                         | 30 Seconds                   |

**Table 2. Availability, DPM and downtime equivalences**

Many systems today strive to achieve five nines availability, which corresponds to less than five minutes of downtime per year.

This percentage availability can also be expressed with the following formula:

$$\text{Percentage of Availability} = (\text{MTBF} * 100) / (\text{MTBF} + \text{MTTR})$$

where MTBF corresponds to the Mean Time Between Failures and MTTR represents the Mean Time To Repair. To reduce MTTR, a router needs to synchronize information among its redundant components. The types of information exchanged between the components could be routing tables, link state information, SNMP, etc. This will allow the router to perform a graceful system recovery during a software maintenance or an outage. Tellabs® 8800 multi-service router accomplishes this with the five nines reliability Tellabs ServiceAssured upgrade.

## SAU Objectives

Typical routers will perform two or three upgrades a year, depending on the maturity of the software or customer-driven enhancements. Typical network downtimes are caused by a change in the state of the network, usually introduced into the network through a software upgrade. For the most part, software upgrades are done to repair or enhance current software operations. Tellabs understands the impact a software upgrade may have on a production network, especially with service providers being held accountable through SLAs. Requiring service providers to reload a node to complete a software upgrade procedure is not an acceptable solution. This is particularly evident with service providers accustomed to the level of availability provided by their Class 5 switches. The objectives of Tellabs ServiceAssured upgrades on the Tellabs® 8800 Multi-service Router Series is to enable service providers to maintain their current software operations with the highest availability-time in the industry. To do so, software upgrades must be done with minimal disruption to customer traffic.

In the past, the operating system of the Tellabs® 8800 series provided non-service affecting software upgrades for the control plane on the Tellabs® 8860 multi-service router and the Tellabs® 8840 multi-service router. Leveraging the Switch Common Control (SCC) redundancy architecture of the Tellabs® 8800 series, service providers can upgrade the software on the standby SCC and use a forced fail-over mechanism to upgrade the active SCC software. All this is done with no impact on customer traffic. While this functionality in itself was superior to methods employed by legacy routers, it did not address the impact of upgrading the software on the Universal Line Cards (ULCs) and the effect on the data-plane traffic. Tellabs® 8800 series with FP4.1 and above provides continuous system availability by restarting critical control plane processes such as routing protocols or signaling stacks automatically during software upgrades. This is achieved through intelligent state checkpointing required to maintain a critical operating state across the platform. With FP4.1 and above, service providers can now perform a complete in-service upgrade on a Tellabs® 8800 series product in a production network.

**Requirements**

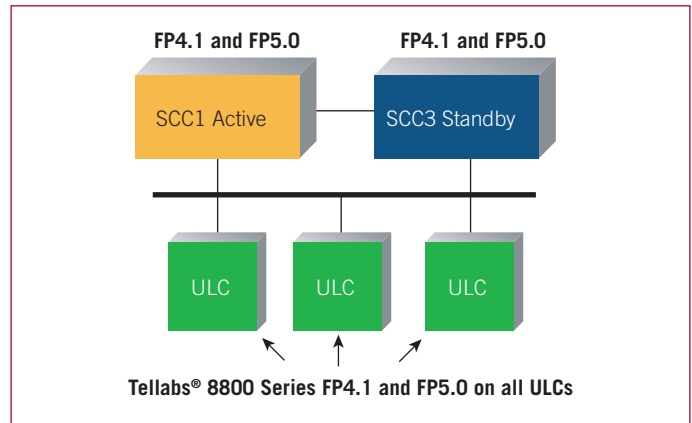
Tellabs ServiceAssured software upgrade is supported on Tellabs® 8800 series platforms that have redundant designs. The upgrade is supported today on the Tellabs® 8860 MSR and Tellabs® 8840 MSR and will be supported with the Tellabs® 8830 MSR in the future. For the Tellabs® 8840 and 8860 MSRs, the minimum configuration would require two SCCs. The minimum software release version required to initiate the Tellabs ServiceAssured upgrade is the Tellabs FP4.1 multi-service operating system. To ensure successful completion of the upgrade, the system should be functional and in a fully redundant state.

**How It Works**

Tellabs ServiceAssured software upgrade is performed via a phased approach. The first phase, referred to as the preparation phase, where the node is taken through a sequence of steps to ensure that it is fully functional and in service at the time of the Tellabs ServiceAssured upgrade. The second phase, the execution phase, is where the node initiates the upgrade process and performs a series of steps to complete the procedure. To ensure minimal user intervention, many of the critical steps in the execution phase are automated.

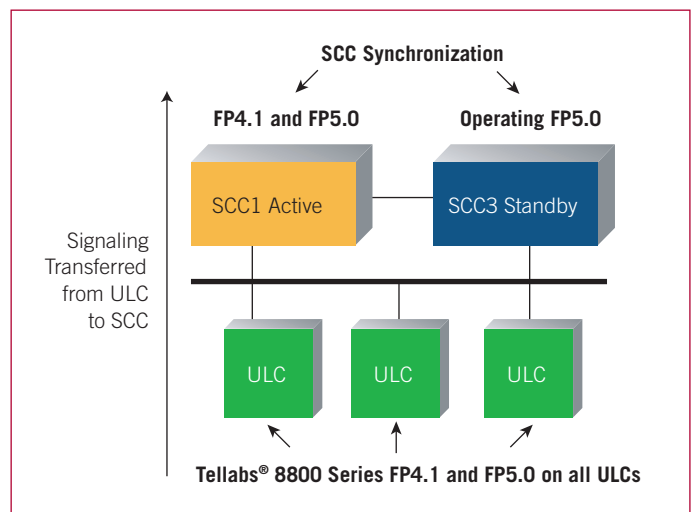
In Figure 1, FP5.0 is downloaded to SCC1 and SCC3 and all three ULCs. Downloading the target software to the SCCs and the ULCs will have no impact on the traffic flow. The SCCs and ULCs can store multiple versions of software. The target software and the current running

software are stored on the Random Access Memory (RAM) and the flash disk. This allows the system to fall back to the original version of code if any failure is encountered during the upgrade procedure. After the download, the system now has FP4.1 as its operating software and FP5.0 as the target software resident in memory.



**Figure 1. Initial System Configuration**

The node is now ready to enter the preparation phase where a sequence of steps verify node redundancy and full operational functionality. Once the preparation phase is completed, the node is ready for the second phase, referred to as the execution phase. The execution phase is completed with minimal user intervention. In the execution phase, the first element that is upgraded is the standby SCC module. This is done by resetting the SCC3 card and having it become operational with FP5.0 software. SCC3 then synchronizes with SCC1 by receiving checkpoint data from SCC1 as shown in Figure 2 below. During this step, data conversions are done and the configuration files are stored on SCC3 with the new FP5.0 database



**Figure 2. SCC3 Upgrade and Synchronization**

format so a complete copy of the database is created and saved in memory on the active and standby SCC. This will enable the operators to revert back to the original mode of operation, if desired. The node is no longer fully redundant, and no configuration changes are allowed. Once the synchronization is completed, the configuration file and the database residing on both SCCs will be identical.

After synchronization, the system is now ready to enter the isolation period. Having the node enter the isolation period ensures that data and configuration integrity is maintained during the execution of the upgrade procedure. During this isolation period, which typically lasts about 2 to 13 minutes (depending on the database size), alarms, routing updates or configuration changes are ignored. The signaling (e.g., RSVP hellos) is transferred from the ULCs to the active SCC for each application. By transferring the signaling to the SCC, the ULC software is no longer in the critical path for maintaining and passing circuit data. IS-IS hellos are also moved to the SCC1 with SCC1 and SCC3 continuously exchanging status information.

The ULCs are then upgraded to FP5.0 through a soft reset as shown in Figure 3 below. By transferring the signaling function to the SCC, we ensure that a soft reset of the ULC will not result in signaling errors and eventual report of a link failure by a remote node. Each ULC reloads the FP5.0 software from RAM. Reloading from RAM minimizes the startup time for the ULCs. ULCs do not write to hardware as traffic continues to be forwarded. All Layer 2 and Layer 3 sessions such as LSPs, OSPF, IS-IS are fully maintained by SCC1 during the reload. The ULCs replay connection information from the active SCC and create a RAM image, but they do not write to the hardware.

After all of the ULCs have completed replay, they rewrite all of the hardware (PP, PM, PS, CAM) at once to match the replayed configuration. The signaling on the SCC is also suspended. The process is synchronized so that connections, microcode, and CAM entries across all ULCs are from the same FP5.0 version before activating the traffic flow again. During this rewrite step, there is about a one-second period where traffic forwarding is halted while the hardware is reinitialized.

Once all the connections, microcode and CAM entries are synchronized across all the ULCs, the signaling functionality is returned to the ULCs.

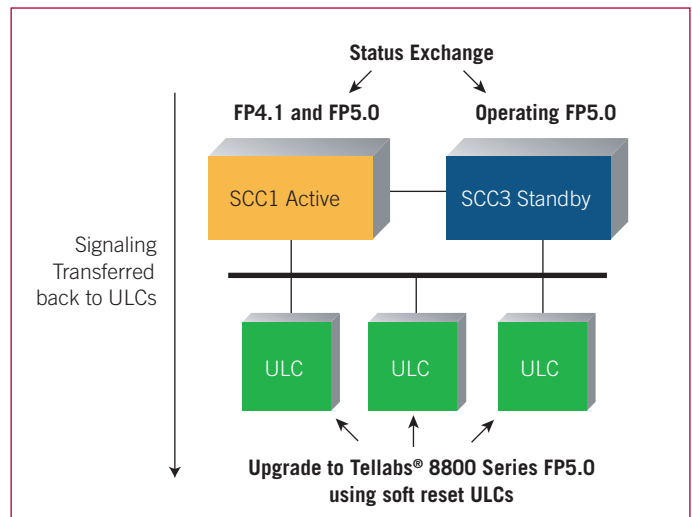


Figure 3. Upgrade of all ULCs to FP5.0

The node now performs an SCC switchover as shown in Figure 4 below, where SCC3 running the target software becomes active and the node isolation is ended. SCC1 is reset and is reloaded with the target software and becomes the new standby after synchronizing with the active SCC. The node is now fully redundant and operational with the target software version.

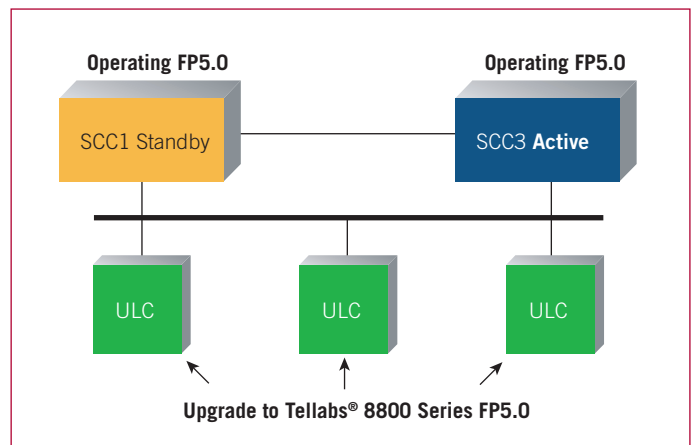


Figure 4. A Fully Upgraded and Redundant System

From a different angle, the entire Tellabs ServiceAssured upgrade process can be shown with the timing diagram in Figure 5 below.

Prior to the preparation phase, the SCCs and ULCs are loaded with the target FP5.0 software. During the preparation phase, the system undertakes internal checks and verifications to ensure that the software upgrade will be successful. The node is checked to make sure that it is fully functional and completely redundant. The operator will then issue an “enable config node upgrade” command that initiates the SCC reset. At this stage, the system is no longer fully redundant. This procedure is identical to previous ways of upgrading non-Tellabs ServiceAssured upgrade aware systems. Once the standby SCC recovers and is operational with FP5.0, the user is then prompted before entering the isolation phase. Prior to entering the isolation phase, any abort or failure reverts the node to the old software (FP4.1) with no effect on network traffic. If the user decides to proceed with the upgrade, the node will enter the isolation mode. During the isolation phase, the signaling is first transferred from the ULCs to the active SCC. The ULCs are then reloaded and finally the SCC performs a

switchover. The hardware on the ULCs is reprogrammed during this period. It is important to note that upon initiating a standby SCC reset until the time the node exits the isolation period, no configuration changes or routing updates are accepted. During the isolation phase, an abort or failure reverts to old software only through a node reset. The purpose of having the isolation mode is to ensure that the upgrade procedure proceeds successfully without any external interruptions. Once the system leaves the isolation phase, the upgrade is completed, and the standby elements are reloaded and synchronized with the active elements. The node at the completion of this stage is fully redundant and operational.

### IS-IS and BGP Graceful Restart

The Tellabs 8800 multi-service router supports the following standards-based graceful restart and fault tolerance mechanisms:

- IS-IS graceful restart
- BGP graceful restart
- LDP fault tolerance
- RSVP stateful redundancy mechanisms

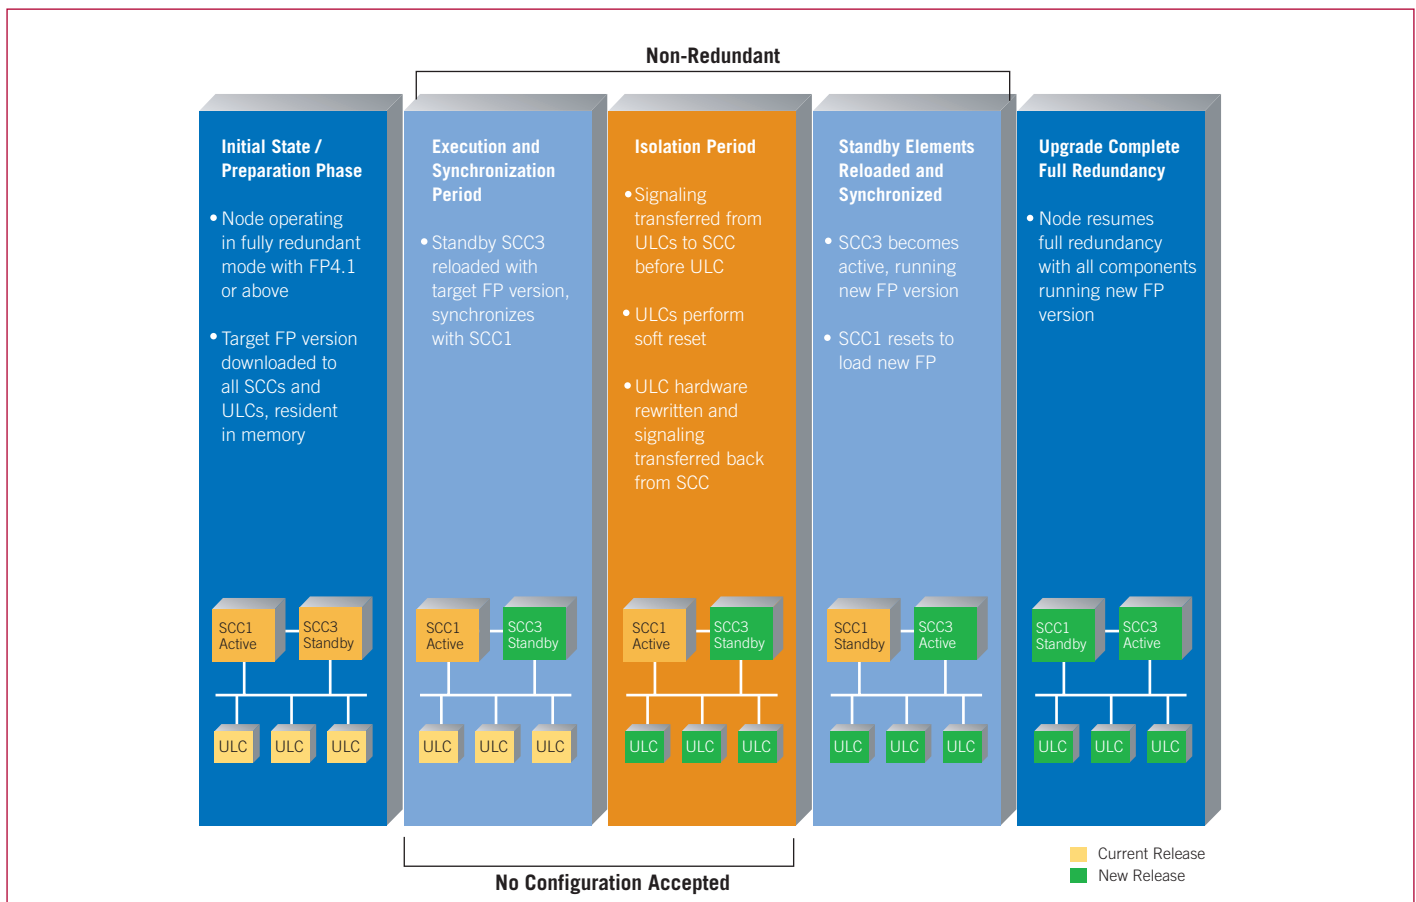


Figure 5. SAU Software Upgrade Timing

A multi-service router from the Tellabs® 8800 Series achieves high availability in IS-IS with the graceful restart mechanism as defined in RFC 3847. When control traffic switches over from the active SCC to the standby SCC, the graceful restart mechanisms are deployed to reacquire the adjacencies and the link state database without any adjacency flaps on the neighboring routers. During the SCC switchover, IS-IS hello PDUs are continuously sent from the ULCs so that the adjacencies are maintained with neighboring routers. IS-IS routes are maintained in the forwarding tables throughout the SCC switchover operation.

For Border Gateway Protocol (BGP), graceful restart is an important feature since this protocol carries a large number of routes relative to other protocols and, consequently, the network may take longer to reach equilibrium after a BGP failure. BGP usually runs at the network edge, affecting the critical link between businesses and the service provider network. Therefore, a failed BGP process can potentially affect multiple networks. In BGP graceful restart, the multi-service router may lose its TCP connection to the peer router. Instead of clearing all of its routes associated with the failed router, the peer router simply marks all routes as stale, and continues to use them to forward packets while waiting for the failed router to re-establish the BGP session. All configuration and various states are actively saved on the standby Switch Controller Card (SCC) of the multi-service router. Throughout the failover process, the IP forwarding table is maintained in the ULC hardware, ensuring that all forwarding of packets is unaffected and BGP peers maintain routes learned from the Tellabs® 8800 series platform.

When the switchover is completed and the new BGP session begins, it will again send BGP capability to its peers. Appropriate flags will be set in the graceful restart capabilities exchange to inform the peer router that the BGP process has restarted.

Other mechanisms employed by the multi-service router include LDP Fault Tolerance as defined by RFC 3479 and RSVP stateful redundancy, all of which work with the redundant components in the Tellabs® 8800 Series to offer the highest availability possible to the end-customers.

## Comparison to Available Industry Solutions

Many of the legacy router products in the industry are not architecturally designed to support a nondisruptive software upgrade process. Whether it's a software crash or a software upgrade, the system will go through a reload process with significant Layer 2 and Layer 3 service disruption. With an average of four to six software upgrades or software crashes a year, the impact on system availability is high if the platform does not support SAU or SAU-like software upgrade. An analysis of routers from other vendors reveals many shortcomings with respect to in-service software upgrade. Most legacy routers today require either an entire node reload or a reload of the system line cards. This results in all the Layer 2 or Layer 3 protocols initializing leading to significant network downtime. Even with redundancy built into the system, depending on the size and complexity of the configuration, the legacy routers have failover times in excess of 90 seconds. When this downtime is factored across a large number of platforms, the impact on the total network availability is so significant that many service providers will avoid system upgrades altogether. Tellabs understands the challenges facing service providers and addresses the critical nature of in-service software upgrades through the Tellabs® ServiceAssured™ feature in the Tellabs® 8800 Series FP4.1 upgrade.

## Summary

During the past few years, network availability has become an increasingly important issue for service providers and their end-customers. Tellabs has responded with the Tellabs® ServiceAssured™ upgrade that helps ensure minimal disruption to existing Layer 2 and Layer 3 customer traffic while upgrading system software. Tellabs® ServiceAssured™ Upgrade provides increased network service availability and protection against planned downtime by offering a true in-service software upgrade procedure to service providers. Deploying Tellabs® ServiceAssured™ upgrade-capable systems at critical network locations will improve system and service availability and will help service providers meet availability conditions set through Service Level Agreements. With Tellabs® ServiceAssured™ upgrades, service providers can offer network service and connectivity at the 99.999% availability level, with no impact to their end-customer's core business. The Tellabs® 8800 series with the Tellabs® ServiceAssured™ upgrade functionality provides the only true carrier-class multi-service platform in the industry.

**FP** — Feature Package  
**MSR** — Multi-Service Router  
**MTBF** — Mean Time Between Failures  
**MTTR** — Mean Time To Repair  
**PLM** — Physical Line Module  
**SAU** — Service Assured Upgrade  
**SCC** — Switch Common Control (Switch Card)  
**SLA** — Service Level Agreement  
**TMOS** — Tellabs Multi-service Operating System  
**ULC** — Universal Line Card

One Tellabs Center  
1415 West Diehl Road  
Naperville, IL 60563 U.S.A.  
Tel: +1 630 798 8800  
Fax: +1 630 798 2000

*The following trademarks and service marks are owned by Tellabs Operations, Inc., or its affiliates in the United States and/or in other countries: TELLABS®, TELLABS and T symbol®, T symbol® and Tellabs® ServiceAssured Upgrade™.*

*All other company names and products mentioned herein may be the property of their respective companies.*

© 2005 Tellabs. All rights reserved.  
74.1538E Rev. B 8/05