Sun StorEdge™ T3 and T3+ Array Field Service Manual

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-4774-10
November 2002, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X / Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, JumpStart, Sun StorEdge, Storage Automated Diagnostic Environment, SunSolve, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 Etats-Unis. Tous droits réservés.

Ce produit ou document est distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X / Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, JumpStart, Sun StorEdge, Storage Automated Diagnostic Environment, SunSolve, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.

Please
Recycle

Regulatory Compliance Statements

Your Sun product is marked to indicate its compliance class:

- Federal Communications Commission (FCC) — USA
- Industry Canada Equipment Standard for Digital Equipment (ICES-003) - Canada
- Voluntary Control Council for Interference (VCCI) — Japan
- Bureau of Standards Metrology and Inspection (BSMI) — Taiwan

Please read the appropriate section that corresponds to the marking on your Sun product before attempting to install the product.

FCC Class A Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note – This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and if it is not installed and used in accordance with the instruction manual, it may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Shielded Cables: Connections between the workstation and peripherals must be made using shielded cables to comply with FCC radio frequency emission limits. Networking connections can be made using unshielded twisted-pair (UTP) cables.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

FCC Class B Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note – This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio / television technician for help.

Shielded Cables: Connections between the workstation and peripherals must be made using shielded cables in order to maintain compliance with FCC radio frequency emission limits. Networking connections can be made using unshielded twisted pair (UTP) cables.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

ICES-003 Class A Notice - Avis NMB-003, Classe A

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

ICES-003 Class B Notice - Avis NMB-003, Classe B

This Class B digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.

VCCI 基準について

クラス A VCCI 基準について

クラス A VCCI の表示があるワークステーションおよびオプション製品は、クラス A 情報技術装置です。これらの製品には、下記の項目が該当します。

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

クラス B VCCI 基準について

クラス B VCCI の表示  があるワークステーションおよびオプション製品は、クラス B 情報技術装置です。これらの製品には、下記の項目が該当します。

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス B 情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。取扱説明書に従って正しい取り扱いをしてください。

BSMI Class A Notice

The following statement is applicable to products shipped to Taiwan and marked as Class A on the product compliance label.

警告使用者：
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻 干擾，在這種情況下，使用者 會被要求採取某些適當的對策。

警告使用者：
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻 干擾，在這種情況下，使用者 會被要求採取某些適當的對策。

Safety Agency Compliance Statements

Read this section before beginning any procedure. The following text provides safety precautions to follow when installing a Sun Microsystems product.

Safety Precautions

For your protection, observe the following safety precautions when setting up your equipment:

- Follow all cautions and instructions marked on the equipment.
- Ensure that the voltage and frequency of your power source match the voltage and frequency inscribed on the equipment's electrical rating label.
- Never push objects of any kind through openings in the equipment. Dangerous voltages may be present. Conductive foreign objects could produce a short circuit that could cause fire, electric shock, or damage to your equipment.

Symbols

The following symbols may appear in this book:

Caution – Caution – There is risk of personal injury and equipment damage. Follow the instructions.

Caution – Caution – Hot surface. Avoid contact. Surfaces are hot and may cause personal injury if touched.

Caution – Caution – Hazardous voltages are present. To reduce the risk of electric shock and danger to personal health, follow the instructions.

Caution – On – Applies AC power to the system.

Depending on the type of power switch your device has, one of the following symbols may be used:

Caution – Off – Removes AC power from the system.

Standby – Standby – The On/Standby switch is in the *standby* position.

Modifications to Equipment

Do not make mechanical or electrical modifications to the equipment. Sun Microsystems is not responsible for regulatory compliance of a modified Sun product.

Placement of a Sun Product

Caution – Caution – Do not block or cover the openings of your Sun product. Never place a Sun product near a radiator or heat register. Failure to follow these guidelines can cause overheating and affect the reliability of your Sun product.

Caution – Caution – The workplace-dependent noise level defined in DIN 45 635 Part 1000 must be 70Db(A) or less.

SELV Compliance

Safety status of I/O connections comply to SELV requirements.

Power Cord Connection

Caution – Caution – Sun products are designed to work with single-phase power systems having a grounded neutral conductor. To reduce the risk of electric shock, do not plug Sun products into any other type of power system. Contact your facilities manager or a qualified electrician if you are not sure what type of power is supplied to your building.

Caution – Caution – Not all power cords have the same current ratings. Household extension cords do not have overload protection and are not meant for use with computer systems. Do not use household extension cords with your Sun product.

Caution – Caution – Your Sun product is shipped with a grounding type (three-wire) power cord. To reduce the risk of electric shock, always plug the cord into a grounded power outlet.

The following caution applies only to devices with a **Standby** power switch:

Caution – Caution – The power switches of this product function as standby type devices only. The power cords serve as the primary disconnect device for the system. ALL power cords must be disconnected to remove power from the product. Be sure to plug the power cords into a grounded power outlet that is nearby the system and is readily accessible.

Lithium Battery

Caution – Caution – On the system control board, there is a lithium battery molded into the real-time clock, SGS No. MK48T59Y, MK48TXXB-XX, MK48T18-XXXPCZ, M48T59W-XXXPCZ, M4T28 XXYYSHZ or MK48T08. Batteries are not customer replaceable parts. They may explode if mishandled. Do not dispose of the battery in fire. Do not disassemble it or attempt to recharge it.

Battery Pack

Caution – Caution - There is a Nickel Metal Hydride battery in the product power supply. Panasonic Model HHR200SCP. There is danger of explosion if the battery is mishandled or incorrectly replaced. Replace only with the same type of Sun Microsystems battery. Do not disassemble it or attempt to recharge it outside the system. Do not dispose of the battery in fire. Dispose of thebattery properly in accordance with local regulations.

System Unit Cover

Caution – Caution – Do not operate Sun products without the top cover in place. Failure to take this precaution may result in personal injury and system damage.

Laser Compliance Notice

Sun products that use laser technology comply with Class 1 laser requirements.

Caution – Caution – Use of controls, adjustments, or the performance of procedures other than those specified herein may result in hazardous radiation exposure.

Einhaltung sicherheitsbehördlicher Vorschriften

Auf dieser Seite werden Sicherheitsrichtlinien beschrieben, die bei der Installation von Sun-Produkten zu beachten sind.

Sicherheitsvorkehrungen

Treffen Sie zu Ihrem eigenen Schutz die folgenden Sicherheitsvorkehrungen, wenn Sie Ihr Gerät installieren:

- Beachten Sie alle auf den Geräten angebrachten Warnhinweise und Anweisungen.
- Vergewissern Sie sich, daß Spannung und Frequenz Ihrer Stromquelle mit der Spannung und Frequenz übereinstimmen, die auf dem Etikett mit den elektrischen Nennwerten des Geräts angegeben sind.
- Stecken Sie auf keinen Fall irgendwelche Gegenstände in Öffnungen in den Geräten. Leitfähige Gegenstände könnten aufgrund der möglicherweise vorliegenden gefährlichen Spannungen einen Kurzschluß verursachen, der einen Brand, Stromschlag oder Geräteschaden herbeiführen kann.

Symbole

Die Symbole in diesem Handbuch haben folgende Bedeutung:

Caution – Achtung – Gefahr von Verletzung und Geräteschaden. Befolgen Sie die Anweisungen.

Caution – Achtung – Hohe Temperatur. Nicht berühren, da Verletzungsgefahr durch heiße Oberfläche besteht.

Caution – Achtung – Gefährliche Spannungen. Anweisungen befolgen, um Stromschläge und Verletzungen zu vermeiden.

Caution – Ein – Setzt das System unter Wechselstrom.

Je nach Netzschaltertyp an Ihrem Gerät kann eines der folgenden Symbole benutzt werden:

Caution – Aus – Unterbricht die Wechselstromzufuhr zum Gerät.

Caution – Wartezustand (Stand-by-Position) - Der Ein-/Wartezustand-Schalter steht auf Wartezustand. Änderungen an Sun-Geräten.

Nehmen Sie keine mechanischen oder elektrischen Änderungen an den Geräten vor. Sun Microsystems übernimmt bei einem Sun-Produkt, das geändert wurde, keine Verantwortung für die Einhaltung behördlicher Vorschriften.

Aufstellung von Sun-Geräten

Caution – Achtung – Um den zuverlässigen Betrieb Ihres Sun-Geräts zu gewährleisten und es vor Überhitzung zu schützen, dürfen die Öffnungen im Gerät nicht blockiert oder verdeckt werden. Sun-Produkte sollten niemals in der Nähe von Heizkörpern oder Heizluftklappen aufgestellt werden.

Caution – Achtung – Der arbeitsplatzbezogene Schalldruckpegel nach DIN 45 635 Teil 1000 beträgt 70Db(A) oder weniger.

Einhaltung der SELV-Richtlinien

Die Sicherung der I/O-Verbindungen entspricht den Anforderungen der SELV-Spezifikation.

Anschluß des Netzkabels

Caution – Achtung – Sun-Produkte sind für den Betrieb an Einphasen-Stromnetzen mit geerdetem Nulleiter vorgesehen. Um die Stromschlaggefahr zu reduzieren, schließen Sie Sun-Produkte nicht an andere Stromquellen an. Ihr Betriebsleiter oder ein qualifizierter Elektriker kann Ihnen die Daten zur Stromversorgung in Ihrem Gebäude geben.

Caution – Achtung – Nicht alle Netzkabel haben die gleichen Nennwerte. Herkömmliche, im Haushalt verwendete Verlängerungskabel besitzen keinen Überlastungsschutz und sind daher für Computersysteme nicht geeignet.

Caution – Achtung – Ihr Sun-Gerät wird mit einem dreiadrigen Netzkabel für geerdete Netzsteckdosen geliefert. Um die Gefahr eines Stromschlags zu reduzieren, schließen Sie das Kabel nur an eine fachgerecht verlegte, geerdete Steckdose an.

Die folgende Warnung gilt nur für Geräte mit Wartezustand-Netzschalter:

Caution – Achtung – Die Ein / Aus-Schalter dieses Geräts schalten nur auf Wartezustand (Stand-By-Modus). Um die Stromzufuhr zum Gerät vollständig zu unterbrechen, müssen Sie die Netzkabel aus der Steckdose ziehen. Alle Netzkabel müssen ausgesteckt sein, um die Stromverbindung zum Produkt zu unterbrechen. Schließen Sie die Stecker der Netzkabel an eine in der Nähe befindliche, frei zugängliche, geerdete Netzsteckdose an.

Lithiumbatterie

Caution – Achtung – Systemsteuerungskarten verfügen über eine Echtzeituhr mit integrierter Lithiumbatterie (Teile-Nr. MK48T59Y, MK48TXXB-XX, MK48T18-XXXPCZ, M48T59W-XXXPCZ, M4T28 XXYYSHZ oder MK48T08). Diese Batterie darf nur von einem qualifizierten Servicetechniker ausgewechselt werden, da sie bei falscher Handhabung explodieren kann. Werfen Sie die Batterie nicht ins Feuer. Versuchen Sie auf keinen Fall, die Batterie auszubauen oder wiederaufzuladen.

Batterien

Caution – Achtung – Das Netzteil des Panasonic-Modells HHR200SCP enthält eine Nickel-Metall-Hydridbatterie. Werden bei der Behandlung oder beim Austausch der Batterie Fehler gemacht, besteht Explosionsgefahr. Tauschen Sie Batterien nur gegen Batterien gleichen Typs von Sun Microsystems aus. Demontieren Sie die Batterie nicht, und versuchen Sie nicht, die Batterie außerhalb des Geräts zu laden. Werfen Sie die Batterie nicht ins Feuer. Entsorgen Sie die Batterie ordnungsgemäß entsprechend den vor Ort geltenden Vorschriften.

Gehäuseabdeckung

Caution – Achtung – Bei Betrieb des Systems ohne obere Abdeckung besteht die Gefahr von Stromschlag und Systemschäden.

Einhaltung der Richtlinien für Laser

Sun-Produkte, die mit Laser-Technologie arbeiten, entsprechen den Anforderungen der Laser Klasse 1.

Caution – Warnung – Die Verwendung von anderen Steuerungen und Einstellungen oder die Durchfhrung von Prozeduren, die von den hier beschriebenen abweichen, knnen gefhrliche Strahlungen zur Folge haben.

Conformité aux normes de sécurité

Ce texte traite des mesures de sécurité qu'il convient de prendre pour l'installation d'un produit Sun Microsystems.

Mesures de sécurité

Pour votre protection, veuillez prendre les précautions suivantes pendant l'installation du matériel :

- Suivre tous les avertissements et toutes les instructions inscrites sur le matériel.
- Vérifier que la tension et la fréquence de la source d'alimentation électrique correspondent à la tension et à la fréquence indiquées sur l'étiquette de classification de l'appareil.
- Ne jamais introduire d'objets quels qu'ils soient dans une des ouvertures de l'appareil. Vous pourriez vous trouver en présence de hautes tensions dangereuses. Tout objet conducteur introduit de la sorte pourrait produire un court-circuit qui entraînerait des flammes, des risques d'électrocution ou des dégâts matériels.

Symboles

Vous trouverez ci-dessous la signification des différents symboles utilisés :

Caution – Attention : risques de blessures corporelles et de dégâts matériels. Veuillez suivre les instructions.

Caution – Attention : surface à température élevée. Evitez le contact. La température des surfaces est élevée et leur contact peut provoquer des blessures corporelles.

Caution – Attention : présence de tensions dangereuses. Pour éviter les risques d'électrocution et de danger pour la santé physique, veuillez suivre les instructions.

Caution – MARCHE : votre système est sous tension (courant alternatif).

Un des symboles suivants sera peut-être utilisé en fonction du type d'interrupteur de votre système:

Caution – ARRET : votre système est hors tension (courant alternatif).

Caution – VEILLEUSE : l'interrupteur Marche / Veilleuse est en position « Veilleuse ».

Modification du matériel

Ne pas apporter de modification mécanique ou électrique au matériel. Sun Microsystems n'est pas responsable de la conformité réglementaire d'un produit Sun qui a été modifié.

Positionnement d'un produit Sun

Caution – Attention : pour assurer le bon fonctionnement de votre produit Sun et pour l'empêcher de surchauffer, il convient de ne pas obstruer ni recouvrir les ouvertures prévues dans l'appareil. Un produit Sun ne doit jamais être placé à proximité d'un radiateur ou d'une source de chaleur.

Caution – Attention : le niveau de pression acoustique au poste de travail s'élève selon la norme DIN 45 635 section 1000, à 70 dB (A) ou moins.

Conformité SELV

Sécurité : les raccordements E / S sont conformes aux normes SELV.

Connexion du cordon d'alimentation

Caution – Attention : les produits Sun sont conçus pour fonctionner avec des alimentations monophasées munies d'un conducteur neutre mis à la terre. Pour écarter les risques d'électrocution, ne pas brancher de produit Sun dans un autre type d'alimentation secteur. En cas de doute quant au type

d'alimentation électrique du local, veuillez vous adresser au directeur de l'exploitation ou à un électricien qualifié.

Caution – Attention : tous les cordons d'alimentation n'ont pas forcément la même puissance nominale en matière de courant. Les rallonges d'usage domestique n'offrent pas de protection contre les surcharges et ne sont pas prévues pour les systèmes d'ordinateurs. Ne pas utiliser de rallonge d'usage domestique avec votre produit Sun.

Caution – Attention : votre produit Sun a été livré équipé d'un cordon d'alimentation à trois fils (avec prise de terre). Pour écarter tout risque d'électrocution, branchez toujours ce cordon dans une prise mise à la terre.

L'avertissement suivant s'applique uniquement aux systèmes équipés d'un interrupteur VEILLEUSE:

Caution – Attention : les commutateurs d'alimentation de ce produit fonctionnent comme des dispositifs de mise en veille uniquement. Ce sont les prises d'alimentation qui servent à mettre le produit hors tension. Vous devez débrancher TOUTES les prises d'alimentation afin de couper l'alimentation du produit. Veillez donc à installer le produit à proximité d'une prise murale facilement accessible.

Batterie au lithium

Caution – Attention : sur la carte de contrôle du système, une batterie au lithium (référence MK48T59Y, MK48TXXB-XX, MK48T18-XXXPCZ, M48T59W-XXXPCZ, M4T28-XXXYYSHZ ou MK48T08) a été moulée dans l'horloge temps réel SGS. Les batteries ne sont pas des pièces remplaçables par le client. Elles risquent d'exploser en cas de mauvais traitement. Ne pas jeter la batterie au feu. Ne pas la démonter ni tenter de la recharger.

Bloc-batterie

Caution – Attention : l'alimentation du produit contient une batterie nickel-hydrure métallique (Panasonic modèle HHR200SCP). Il existe un risque d'explosion si cette batterie est manipulée de façon erronée ou mal mise en place. Ne remplacez cette batterie que par une batterie Sun Microsystems du même type. Ne la démontez pas et n'essayez pas de la recharger hors du système. Ne faites pas brûler la batterie mais mettez-la au rebut conformément aux réglementations locales en vigueur.

Couvercle

Caution – Attention : il est dangereux de faire fonctionner un produit Sun sans le couvercle en place. Si l'on néglige cette précaution, on encourt des risques de blessures corporelles et de dégâts matériels.

Conformité aux certifications Laser

Les produits Sun qui font appel aux technologies lasers sont conformes aux normes de la classe 1 en la matière.

Caution – Attention : l'utilisation de contrôles, de réglages ou de performances de procédures autre que celle spécifiée dans le présent document peut provoquer une exposition à des radiations dangereuses.

Normativas de seguridad

El siguiente texto incluye las medidas de seguridad que se deben seguir cuando se instale algún producto de Sun Microsystems.

Precauciones de seguridad

Para su protección observe las siguientes medidas de seguridad cuando manipule su equipo:

- Siga todos los avisos e instrucciones que se indican en el equipo.
- Asegúrese de que el voltaje y la frecuencia de la red eléctrica concuerdan con las descritas en las etiquetas de especificaciones eléctricas del equipo.
- No introduzca nunca objetos de ningún tipo a través de los orificios del equipo. El voltaje puede ser peligroso. Los objetos extraños conductores de la electricidad pueden producir cortocircuitos que provoquen un incendio, descargas eléctricas o daños en el equipo.

Símbolos

En este libro aparecen los siguientes símbolos:

Caution – Precaución – Existe el riesgo de lesiones personales y daños al equipo. Siga las instrucciones.

Caution – Precaución – Superficie caliente. Evite el contacto. Las superficies están calientes y pueden causar daños personales si se tocan.

Caution – Precaución – Voltaje peligroso presente. Para reducir el riesgo de descarga y daños para la salud siga las instrucciones.

Caution – Encendido – Aplica la alimentación de CA al sistema.

Según el tipo de interruptor de encendido que su equipo tenga, es posible que se utilice uno de los siguientes símbolos:

Caution – Apagado – Elimina la alimentación de CA del sistema.

Caution – En espera – El interruptor de Encendido/En espera se ha colocado en la posición de *En espera.*

Modificaciones en el equipo

No realice modificaciones de tipo mecánico o eléctrico en el equipo. Sun Microsystems no se hace responsable del cumplimiento de las normativas de seguridad en los equipos Sun modificados.

Ubicación de un producto Sun

Caution – Precaución – Para asegurar la fiabilidad de funcionamiento de su producto Sun y para protegerlo de sobrecalentamientos no deben obstruirse o taparse las rejillas del equipo. Los productos Sun nunca deben situarse cerca de radiadores o de fuentes de calor.

Caution – Precaución – De acuerdo con la norma DIN 45 635, sección 1000, se admite un nivel de presión acústica para puestos de trabajo máximo de 70Db(A).

Cumplimiento de la normativa SELV

El estado de la seguridad de las conexiones de entrada/salida cumple los requisitos de la normativa SELV.

Conexión del cable de alimentación eléctrica

Caution – Precaución – Los productos Sun están diseñados para trabajar en una red eléctrica monofásica con toma de tierra. Para reducir el riesgo de descarga eléctrica, no conecte los productos Sun a otro tipo de sistema de alimentación eléctrica. Póngase en contacto con el responsable de mantenimiento

o con un electricista cualificado si no está seguro del sistema de alimentación eléctrica que existe en su edificio.

Caution – Precaución – No todos los cables de alimentación eléctrica tienen la misma capacidad. Los cables de tipo doméstico no están provistos de protecciones contra sobrecargas y por tanto no son apropiados para su uso con computadores. No utilice alargadores de tipo doméstico para conectar sus productos Sun.

Caution – Precaución – Con el producto Sun se proporciona un cable de alimentación con toma de tierra. Para reducir el riesgo de descargas eléctricas conéctelo siempre a un enchufe con toma de tierra.

La siguiente advertencia se aplica solamente a equipos con un interruptor de encendido que tenga una posición "En espera":

Caution – Precaución – El interruptor de encendido de este producto funciona exclusivamente como un dispositivo de puesta en espera. Los enchufes de la fuente de alimentación están diseñados para ser el elemento primario de desconexión del equipo. Debe desconectar TODOS los enchufes de alimentación del equipo antes de desconectar la alimentación. El equipo debe instalarse cerca del enchufe de forma que este último pueda ser fácil y rápidamente accesible.

Batería de litio

Caution – Precaución – En las placas de control del sistema hay una batería de litio insertada en el reloj de tiempo real, tipo SGS Núm. MK48T59Y, MK48TXXB-XX, MK48T18-XXXPCZ, M48T59W-XXXPCZ, M4T28-XXYYSHZ o MK48T08. El usuario no debe reemplazar las baterías por sí mismo. Pueden explotar si se manipulan de forma errónea. No arroje las baterías al fuego. No las abra o intente recargarlas.

Paquete de pilas

Caution – Precaución – Existe una pila de hidruro metálico de níquel en el sistema de alimentación de la unidad Panasonic modelo HHR200SCP. Existe riesgo de estallido si el paquete de pilas se maneja sin cuidado o se sustituye de manera indebida. Las pilas sólo deben sustituirse por el mismo tipo de pilas de Sun Microsystems. No las desmonte ni intente recargarlas fuera del sistema. No arroje las pilas al fuego. Deséchelas siguiendo el método indicado por las disposiciones vigentes.

Tapa de la unidad del sistema

Caution – Precaución – Es peligroso hacer funcionar los productos Sun sin la tapa superior colocada. El hecho de no tener en cuenta esta precaución puede ocasionar daños personales o perjudicar el funcionamiento del equipo.

Aviso de cumplimiento con requisitos de láser

Los productos Sun que utilizan la tecnología de láser cumplen con los requisitos de láser de Clase 1.

Caution – Precaución – El manejo de los controles, los ajustes o la ejecución de procedimientos distintos a los aquí especificados pueden exponer al usuario a radiaciones peligrosas.

GOST-R Certification Mark

Nordic Lithium Battery Cautions

Norge

Caution – A D V A R S E L – Litiumbatteri — Eksplosjonsfare.
Ved utskifting benyttes kun batteri som anbefalt av apparatfabrikanten. Brukt batteri returneres apparatleverandøren.

Sverige

Caution – VARNING – Explosionsfara vid felaktigt batteribyte. Använd samma batterityp eller en ekvivalent typ som rekommenderas av apparattillverkaren. Kassera använt batteri enligt fabrikantens instruktion.

Danmark

Caution – ADVARSEL! – Litiumbatteri — Eksplosionsfare ved fejlagtig håndtering. Udskiftning må kun ske med batteri af samme fabrikat og type. Levér det brugte batteri tilbage til leverandøren.

Suomi

Caution – VAROITUS – Paristo voi räjähtää, jos se on virheellisesti asennettu. Vaihda paristo ainoastaan laitevalmistajan suosittelemaan tyyppiin. Hävitä käytetty paristo valmistajan ohjeiden mukaisesti.

Contents

Figures

Tables

Preface

The *Sun StorEdge T3+ Array Field Service Manual* is designed to provide the qualified service-trained maintenance provider with sufficient information to effectively troubleshoot and resolve any Sun StorEdge™ T3+ array failure. The procedures in this manual describe how to isolate the failure, remove and replace component(s), effectively reconfigure the module and system, and place the product back into the customer's network.

Before You Read This Book

Make sure you have prepared by reviewing the *Sun StorEdge T3+ Array Installation and Configuration Manual*; *Sun StorEdge T3+ Array Administrator's Manual;* and *Sun StorEdge T3+ Array Release Notes*. Work with the site system administrator to determine if any external hardware or software products are required to repair this device.

How This Book Is Organized

This manual is organized as follows:

Chapter 1 provides a troubleshooting overview on the Sun StorEdge T3+ array.

Chapter 2 describes how to connect to and boot the Sun StorEdge T3+ array.

Chapter 3 provides the qualified service provider with troubleshooting techniques for the Sun StorEdge T3+ array.

Chapter 4 describes how to monitor and replace the controller card, and upgrade the firmware.

Chapter 5 describes how to monitor and replace the disk drives, and upgrade the firmware.

Chapter 6 describes how to monitor and replace the interconnect card, and upgrade the firmware.

Chapter 7 describes how to replace the power and cooling unit and monitor the UPS.

Chapter 8 describes how to diagnose and correct back-end FC-AL drive loop problems with the Sun StorEdge T3+ array.

Chapter 9 describes how to replace the chassis/backplane assembly.

Chapter 10 describes how to reconfigure the Sun StorEdge T3+ array into partner groups and single controller units.

Appendix A contains part numbers and illustrations of field-replaceable units.

Appendix B lists the Sun StorEdge T3+ array defaults.

Appendix C contains a description of the messages that can be reported by the array.

Appendix D contains descriptions of the commands supported by the Sun StorEdge T3+ array.

Appendix E lists the FC-AL loop identified by AL_PA switch and setting values.

Appendix F contains a blank worksheet for the qualified service provider to make notes at each customer site.

Using UNIX Commands

This document contains some information on basic UNIX® commands and procedures such as booting the devices. For more information outside of this document, see the following:

- AnswerBook2™ online documentation for the Solaris™ software environment
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1 Typographic Conventions

Typeface	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `% You have mail.`
`AaBbCc123`	What you type, when contrasted with on-screen computer output	`% su` `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized, glossary terms	Read Chapter 6 in the *User's Guide.* These are called *class* options. The user *must* be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type `rm` *filename.*

Shell Prompts

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	*machine_name*%
C shell superuser	*machine_name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Latest array updates	*Sun StorEdge T3+ Array Release Notes*	816-4771
Installation overview	*Sun StorEdge T3+ Array Start Here*	816-4768
Safety procedures	*Sun StorEdgeT3+ Array Regulatory and Safety Compliance Manual*	816-0774
Site preparation	*Sun StorEdge T3+ Array Site Preparation Guide*	816-0778
Installation and Configuration	*Sun StorEdge T3+ Array Installation and Configuration Manual*	816-4769
Administration	*Sun StorEdge T3+ Array Administrator's Manual*	816-4770
Cabinet installation	*Sun StorEdge T3+ Array Installation and Configuration Manual*	806-7979
Disk drive specifications	*18 Gbyte, 1-inch, 10K rpm Disk Drive Specifications*	806-1493
	36 Gbyte, 10K rpm 1-Inch Disk Drive Specifications	806-6383
	73 Gbyte, 10K rpm 1.6-Inch Disk Drive Specifications	806-4800
Host Bus Adapters	*Sun Enterprise 6x00/5x00/4x00/3x00 Systems SBus+ and Graphics+ I/O Boards Installation Guide*	805-2704
	Sun StorEdge PCI FC-100 Host Adapter Installation	805-3682
	Sun StorEdge SBus FC-100 Host Adapter Installation and Service Manual	806-7532
	Sun StorEdge PCI Single Fibre Channel Network Adapter Installation Guide	806-7532
	Sun StorEdge PCI Dual Fibre Channel Host Adapter Installation Guide	806-4199

Application	Title	Part Number
	Sun StorEdge Compact PCI Dual Fibre Channel Network Adapter Installation and User's Guide	816-0241
Testing the array	*Storage Automated Diagnostic Environment User's Guide*	816-3142
	Storage Automated Diagnostic Environment Version 2.0.06.010 Release Notes[1]	816-3141

1. Can be found at http://webhome.central/storade/.

Accessing Sun Documentation Online

You can access a select group of Sun™ technical documentation on the Web. You can browse the documentation archive at:

`http://`www.sun.com/products-n-solutions/hardware/docs

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

`docfeedback@sun.com`

Please include the part number (816-4774-10) of your document in the subject line of your email.

CHAPTER 1

Troubleshooting Overview

This chapter provides an introduction to some of the tools available to troubleshoot the Sun StorEdge T3+ array and describes the following sections:

Network Storage Overview

An understanding of a *network storage* environment is required before any troubleshooting can take place.

Each Sun StorEdge T3+ array can be configured with a maximum of two RAID volumes. If desired, these volumes can be partitioned into up to 16 distinct logical unit numbers (LUNs). Volumes need not be partitioned with an equal number of LUNs, however, the total must not exceed 16.

The Sun StorEdge T3+ array partner group consists of two independent controller RAID units sharing only one of the controllers for system management. If one controller fails, the system management facilities fail over to the other RAID controller. This configuration gives the partner group redundancy.

Thus, when configured as a Sun StorEdge T3+ array enterprise configuration, which consists of two arrays, a maximum of four RAID volumes (LUNs) are available to the server for data delivery and retrieval. Although an additional two volumes can be added to the two available in a workgroup configuration, the maximum total number of LUNs remains at 16.

The interruption of data can happen anywhere on the storage network. This manual addresses data interruption problems from the output of the host to the Sun StorEdge T3+ array, and to the individual components in the array.

Maintenance Precaution

After configuring a system, always record the following data to prepare for the possibility of having to perform a recovery procedure:

- Array block size
- Multipathing settings
- Volume configuration
- Volume slicing configuration
- LUN masking settings

Error Messages and Logs

Both the Sun StorEdge T3+ array and the host server create log message files of system conditions and events. These log files are the most useful *immediate* tools for troubleshooting.

Sun StorEdge T3+ Array Generated Messages

A `syslog` daemon in the Sun StorEdge T3+ array writes system error message logs to a location determined by the site system administrator. Consult with the site system administrator to obtain access to this log. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for instructions on setting up remote logging.

Host-Generated Message

A `syslog` daemon in the host hardware writes system error message logs to `/var/adm/messages`.

The data host sees an array or enterprise configuration as a group of up to 16 LUNs. A host-generated error message could indicate that the host cannot communicate with the array through the Fibre Channel-Arbitrated Loop (FC-AL) channel, or that an excessive number of channel errors are occurring. If the host loses access to the array through the channel connection, then any host messages regarding the array will refer only to the LUNs.

In a partner group configuration where multi-pathing failover has been established, the failure of a channel path or array controller causes the host to redirect I/O from the failed channel to the second FC-AL connection.

A variety of software logging tools monitor the various branches of the storage network. When an error is detected, the error's severity level is categorized and classified. Errors are reported or logged according to severity level (TABLE 1-1).

TABLE 1-1 Levels of Message Notification

Message Level	Description
Error	Indicates a critical system or storage network event or failure, requiring immediate intervention or attention
Warning	Indicates a possible system or storage network event or failure, requiring eventual intervention
Notice	Indicates a system event that could be a normal periodic notification, a system fault, operator keyboard commands, or a result of other events
Information	Indicates a system event that has no impact upon the system or storage networks ability to perform tasks.

The syntax of the error message uses a field-replaceable unit (FRU) identifier to refer to a particular FRU in a Sun StorEdge T3+ array (TABLE 1-2). This identifier contains a unit constant (u), the unit number (*n*), a FRU constant (`ctr` for controller card, `pcu` for power and cooling unit, `l` for unit interconnect card, `d` for disk drive), and the FRU number (*n*).

TABLE 1-2 FRU Identifiers

FRU	Identifier	Unit number
Controller card	u*n*`ctr`	*n* = unit number (1, 2,...)

TABLE 1-2 FRU Identifiers *(Continued)*

FRU	Identifier	Unit number
Power and cooling unit	u*n*pcu*n*	*n* = unit number (1, 2,...) *n* = pcu number (1, 2)
Unit interconnect card	u*n*l*n*	*n* = unit number (1, 2,...) *n* = interconnect number (1, 2)
Disk drive	u*n*d*n*	*n* = unit number (1, 2,...) *n* = disk drive number (1 to 9)

Sun Storage Automated Diagnostic Environment

The Storage Automated Diagnostic Environment is a host-based online health and diagnostic monitoring tool for storage area network (SAN) and direct-attached storage (DAS) devices. It can be configured to monitor on a 24-hour basis, collecting information that enhances the reliability, availability, and serviceability (RAS) of the storage devices.

The Storage Automated Diagnostic Environment offers the following features:

- A common web-based user interface for device monitoring and diagnostics
- Distributed test invocation by means of lists or topology
- Topology grouping for multi-level hosts and components
- Alternate master support for redundancy
- Revision checking
- Remote notification through SRS, SRS/NetConnect, RSS, HTTP, SSTR, and SMTP Providers, or email
- Support for storage area networks (SANs) and direct-attached storage (DAS) devices

The Storage Automated Diagnostic Environment can be downloaded from the Sun web site. See the *Storage Automated Diagnostic Environment User's Guide* for instructions.

Errors in the host data channel are outside of the scope of the *Sun StorEdge T3+ Array Field Service Manual*. Host-to-array channel failures occur when the connection between the array and the host is either severed or intermittent. The components that make up this data channel connection can include:

- Host bus adapter (HBA), which resides on the host

- Gigabit interface converter (GBIC) adapter, used to connect the FC-AL cable to an SBus HBA
- Fibre Channel cable that connects the array to the host
- Media interface adapter (MIA), which converts the light source from the host to an electron source for use in the array
- Channel interface port in the array
- Fibre Channel switches connecting the host to the storage in a SAN

To determine failures in the data path, use a host-based application diagnostics product, such as the Sun Storage Automated Diagnostic Environment, for the Solaris operating environment.

Static Electricity Precautions

Follow these procedures to prevent damaging the FRUs:

- Remove plastic, vinyl, and foam from the work area.
- Before handling a FRU, discharge any static electric charge by touching a ground surface.
- Wear an antistatic wrist strip.
- Do not remove a FRU from its antistatic protective bag until you are ready to install it.
- When removing a FRU from the array, immediately place it in an antistatic bag and packaging.
- Handle a FRU only by its edges and avoid touching the circuitry.
- Do not slide a FRU over any surface.
- Limit body movement (which builds up static electricity) during FRU installation.

CHAPTER 2

Connecting to the Sun StorEdge T3+ Array

This chapter describes how to connect to the Sun StorEdge T3+ array and contains the following sections:

Establishing a Serial Port Connection

The serial port is a direct connection to the array from any serial port on any host or system. Individual commands can be run to query and repair the unit from this interface using the command-line interface (CLI). The serial port connection provides the following advantages over the Telnet connection:

- Boot messages are displayed when the array boots.
- The `tftp boot` configuration is available.
- EPROM access is available.
- Useful for debugging RARP/IP address assignment issues.
- Array-specific troubleshooting commands can be issued to each controller in an enterprise configuration.

The status of the array unit can quickly be determined from the CLI. The `syslog` file of the array file system contains a record of events that have occurred in the unit.

To start a serial connection and session with the array:

1. **Connect a serial cable from the serial port on the array master unit to any host system available serial port.**

Note – The serial cables used by the Sun StorEdge T3+ arrays are different. Both cables are supplied in the F370-4119-02 Diagnostic Kit. The T3 array uses the serial cable with RJ-11 connectors and the T3+ array uses the serial cable with RJ-45 connectors.

The serial port on the array is on the controller card backplane.

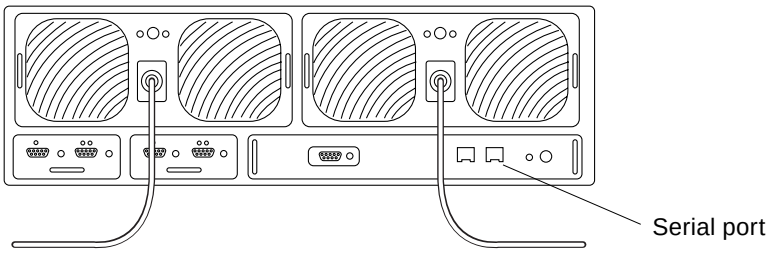

FIGURE 2-1 Serial Port Location

2. **On the host system, open a terminal window, type** `tip`**, the baud rate, and the serial port designation.**

For example:

```
mymachine# tip -9600 /dev/ttya
connected
Password:
Invalid name.
Login:root
Password:

T3B Release 2.00 2001/04/02 15:21:29 (192.168.209.243)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.

:/:<1>
```

If the Sun StorEdge T3 array is being booted, the following message is displayed:

```
auto boot is enabled
hit the RETURN key within 3 seconds to cancel...
```

In a boot situation, if the Return key is pressed within 3 seconds, the array stops booting and the EPROM takes control of the array. If the Return key is not pressed, the array continues to boot. Note that in a partner group the alternate master unit continues to boot and appears as the master if the boot sequence of the master is stopped.

The following commands are available for use at the EPROM level:

- `boot`
- `reset`
- `set`
- `id`

Once the array has fully booted, all the commands available through the CLI are accessible.

Note – Use the login prompt to set the IP address, netmask, and hostname instead of using the EP prompt. Setting these parameters at the EPROM level will be lost.

For more information on serial connections, see:

- *Sun StorEdge T3+ Array Administrator's Manual* for instructions on setting up remote logging
- "Checking Array Boot Status" on page 27

Establishing a Telnet Session

The Telnet session is a direct network link to the array unit through the command-line interface (CLI). You can execute individual commands to query and repair the unit from this interface. The Telnet session requires access to the unit's Ethernet network. The CLI can then be run from any host that can access the array subnetwork.

The advantages that a Telnet connection provides over a serial port connection are as follows:

- You can have multiple windows open for each array.

- The Telnet connection provides a faster interface than the serial port connection, which can be useful for displaying `syslog` information.

You can quickly determine the status of the array unit from the CLI. The `syslog` file on the array file system contains a record of events that have occurred in the unit and can also be examined through the CLI.

To open a Telnet connection and start a session with the array:

1. **On the management host, use the `telnet` command with the array name (or IP address) to connect to the array.**

   For example, to telnet to a array named T3-1:

```
mgmt-host# telnet T3-1
Trying 123.123.123.1...
Connected to T3-1.
Escape character is '^]'.

Telnet session (123.123.123.1)

Login: root
Password: passwd

T3B Release 2.1 2002/04/02 15:21:29 (192.168.209.243)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.

:/:<1>
```

   where *password* is the root password.

2. **Verify the array has a root password by typing it at the prompt.**

   If no root password is set on the system, press Return at the password prompt to enter the CLI. Use the `password` command to establish a password.

3. **To view the available commands on the array, type `help` at the prompt.**

```
:/:<1>help
arp       cat       cd        cmp     cp      date      echo   head
help      ls        mkdir     mv      ping    pwd       rm     rmdir
tail      touch
boot      disable   disk      enable  fru     id        logger lpc
more      ntp       passwd    port    proc    reset     set    shutdown
sync      sys       tzset     ver     vol     volslice  ep
refresh   route     ofdg      lun     hwwn
```

For more information on how to set up the `syslog` file and interpret it, refer to the *Sun StorEdge T3+ Array Administrator's Manual* for instructions on setting up remote logging. For information on how to use the CLI commands, see *Sun StorEdge T3+ Array Administrator's Manual*.

Establishing an FTP Session

To establish an FTP session:

1. **Start an FTP session from the management host to the array.**

   For example:

```
mgmt-host:/:<15>ftp 123.123.123.2
Connected to 123.123.123.2.
Escape character is '^]'.

Telnet session 123.123.123.2 (NUPPC/2.0.0-G) ready

Name (123.123.123.2:root):
```

2. **Log in to the array by typing** `root`**.**

```
Name (123.123.123.2:root): root

331 Password required for root.
Password: password
230 User root logged in.
ftp>
```

where *password* is the root password.

Note – Be sure to set the Binary mode if transferring firmware.

Note – If the root password has not been set, the FTP login to the array will fail.

Using `tftpboot` to Boot a Single Array or a Partner Group Remotely

If you have a partner group that cannot boot on its own, you can use `tftp boot` to boot it remotely.

Note – The tftpboot server must be on the same subnet as the array.

To remotely boot a Sun StorEdge T3+ array:

1. **Set up the remote server.**

   See “Configuring a Server for Remote Booting” on page 16.

2. **Unplug the Ethernet cable connected to the alternate master.**

   Leave the Ethernet cable on the master connected.

3. **Get to the array EPROM as described in “Establishing a Serial Port Connection” on page 7.**

4. **Set the array boot mode to `tftpboot`.**

```
T3-1>set bootmode tftp
T3-1>set
bootmode  auto
bootdelay 3
sn        112035
ip        10.4.35.134
netmask   255.255.255.0
gateway   10.4.35.1
tftphost  123.123.123.6
tftpfile  releases/nb210/nb210p20.bin
hostname  qatest
timezone  GMT+00
vendor    0301
model     501-5710-02(51)
revision  020100
logto     /Aug9
loglevel  3
rarp      off
mac       00:20:f2:00:03:b9
```

5. **Set** `tftphost` **IP address and** `tftp` **filename.**

```
T300-EP>set tftphost 123.123.123.6
T300-EP>set tftpfile filename.bin
T300-EP>set
bootmode    tftp
bootdelay   3
sn          000596
ip          123.123.123.99
netmask     255.255.255.0
gateway     129.153.49.254
tftphost    129.153.49.2
tftpfile    nb210.bin
hostname    purple31
timezone
vendor      0301
model       501-5710-02(51)
revision    0200
logto       /syslog (* in nvram)
loglevel    4
rarp        on
mac         00:20:f2:00:93:24
```

6. **Reset the master to initiate the** `tftp` **boot cycle.**

```
T3-1>reset
Starting...

T3-1 Release 2.10 1999/11/24 13:05:57 (123.123.123.3)
Copyright (C) 1997-1999 Sun Microsystems, Inc.
All Rights Reserved.

Found units: [u1-ctr]
tftp boot is enabled
hit the RETURN key within 3 seconds to cancel...
Initializing TFTP...
Loading 123.123.123.6:nb113.bin
...
...
login:
```

7. **Copy the firmware from the** `tftp` **boot server to the Sun StorEdge T3+ array.**

Note that the `ftp` command is initiated from the TFTP server since at this point you are no longer on the array.

```
mgmt_host# ftp 123.123.123.3
Connected to 123.123.123.3.
220 123.123.123.3 pSOSystem FTP server (NUPPC/2.0.0-G) ready.
Name (123.123.123.3:root): root
331 Password required for root.
Password:
230 User root logged in.
ftp> lcd /tftpboot
Local directory now /tftpboot
ftp> bin
200 Type set to I.
ftp> put filename.bin
200 PORT command successful.
150 Opening BINARY mode data connection for filename.bin.
226 Transfer complete.
local: filename.bin remote: filename.bin
2514468 bytes sent in 51 seconds (47.87 Kbytes/s)
ftp>
```

Where *filename*`.bin` is the name of the current firmware file. For example, `nb113.bin`.

8. **Boot the newly transferred controller firmware image on the master.**

```
:/:<3>boot -i filename.bin
file header: size 265e14, checksum be4ec46, start 20010, base 20000
```

This copies the firmware to the bootable reserved areas on the local disk.

9. **Set the bootmode back to auto**

(If you forget this step, the system will continue doing `tftpboots`.)

```
:/:<4>set bootmode auto
```

10. **Reset the system.**

```
:/:<5>reset
Reset the system, are you sure? [N]: y
```

11. **Reconnect the Ethernet cable to the alternate master.**

Configuring a Server for Remote Booting

If a Sun StorEdge T3+ array is unable to boot you can use `tftboot` to reload the firmware. This requires configuring a remote server.

To configure a remote server to `tftp` boot a Sun StorEdge T3+ array, follow these steps:

1. **In a user file system, create a directory on the server called** `tftpboot`**.**

```
boothost# mkdir /tftpboot
```

2. **Set permissions to allow users read/write access.**

```
boothost# chmod 777 /tftpboot
```

3. **Copy the Sun StorEdge T3+ array boot code into the** `tftpboot` **directory.**

```
boothost# cp nbnnn.bin /tftpboot
```

Where `nb`*nnn*`.bin` is the current boot-code file identification number. For example, `nb101.bin`.

4. **Verify that** `/tftpboot/nb`*nnn*`.bin` **is readable.**

```
boothost# chmod 755 /tftpboot/nbnnn.bin
```

5. **Edit the** `/etc/inetd.conf` **file and uncomment the** `tftp` **line.**

```
tftp dgram udp wait root /usr/sbin/in.tftpd in.tftpd -s /tftpboot
```

6. **Restart** `inetd`**.**

```
boothost# ps -eaf | grep inetd
root 140 1 0 Feb 08 ? 0:00 /usr/sbin/inetd -s
root 7715 7701 0 11:22:32 pts/18 0:00 grep inetd
# kill -HUP 140
```

Setting Up Remote Logging

The Sun StorEdge T3+ array can provide remote notification of array events to designated hosts using Simple Network Management Protocol (SNMP) traps. To enable SNMP notification, edit the `/etc/syslog.conf` and the `/etc/hosts` files on the array to configure system message logging. Because files cannot be edited on the array, `ftp` the files to a host to make the edits and then `ftp` the files back to the array.

Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for instructions on setting up remote logging.

CHAPTER 3

Diagnosing T3+ Array Problems

This chapter provides the qualified service provider with troubleshooting techniques for the Sun StorEdge T3+ array and contains the following sections:

Diagnostic Information Sources

TABLE 3-1 summarizes the diagnostic tools available to you.

TABLE 3-1 Diagnostic Functions and Tools

Function	Tools That Can Be Used
Array boot monitoring	LEDs, CLI(S),
Array boot PROM commands	CLI(S)
Host data path diagnosis	SADE
Internal monitoring	LEDs, CLI (E), CLI(S), SNMP, SADE, syslog, SRS

TABLE 3-1 Diagnostic Functions and Tools *(Continued)*

Function	Tools That Can Be Used
Configuration	LEDs, CLI (E), CLI(S)
System (admin domain) configuration	CLI(E), CLI(S)
System (admin domain) monitoring	CLI(E), CLI(S), SRS, SNMP(CA), syslog(CA)
Version level check	CLI(E), CLI(S)
LUN configuration	CM, CLI(E), CLI(S)
FRU failure monitoring	LED, CLI(E), CLI(S), SRS, SNMP(CA), syslog(CA)
Performance monitoring	CLI(E), CLI(S), SNMP(CA), syslog(CA)
Firmware download	CLI(E)
Syslog access (mgmt host)	CLI(E), CLI(S), syslog(CA and SADE with 2nd copy of SADE running on management host with ethernet connection to array
Loop resiliency check (manual)	OFDG, CLI(E), CLI(S)
Manual loop resiliency check	OFDG, CLI(E), CLI(S)
Clear supervisor password	CLI(S)
Host data path diagnosis	SADE
Statistics logging	syslog(CA) and SADE with 2nd copy of SADE application running on management host with Ethernet connection to array
Service commands	CLI(E), CLI(S)
Mfg / repair commands	CLI(E), CLI(S)

- LED = Light emitting diodes on the array.
- CLI(E) = Command-line utilities run via ethernet connection as described in *Sun StorEdge T3+ Array Administrator's Manual.*
- CLI(S) = Command-line utilities run via a serial connection as described in "Establishing a Serial Port Connection" on page 7.
- OFDG = Off-line Drive Diagnostic utility as described in "Using the `ofdg` Diagnostic Utility" on page 111.
- SNMP = Simple Network Monitoring Protocol as described in *Sun StorEdge T3+ Array Administrator's Manual.*
- SNMP(CA) = Simple Network Monitoring Protocol used with a customer-written application as described in *Sun StorEdge T3+ Array Administrator's Manual.*

- SADE = The Storage Automated Diagnostic Environment application as described in *Storage Automated Diagnostic Environment User's Guide*.
- syslog = Sun StorEdge T3+ array `syslog` file.
- syslog(CA) = Sun StorEdge T3+ array `syslog` with customer-written application.
- SRS = Sun Remote Service.

Troubleshooting Flow Charts

The following three charts illustrate typical diagnostic procedures.

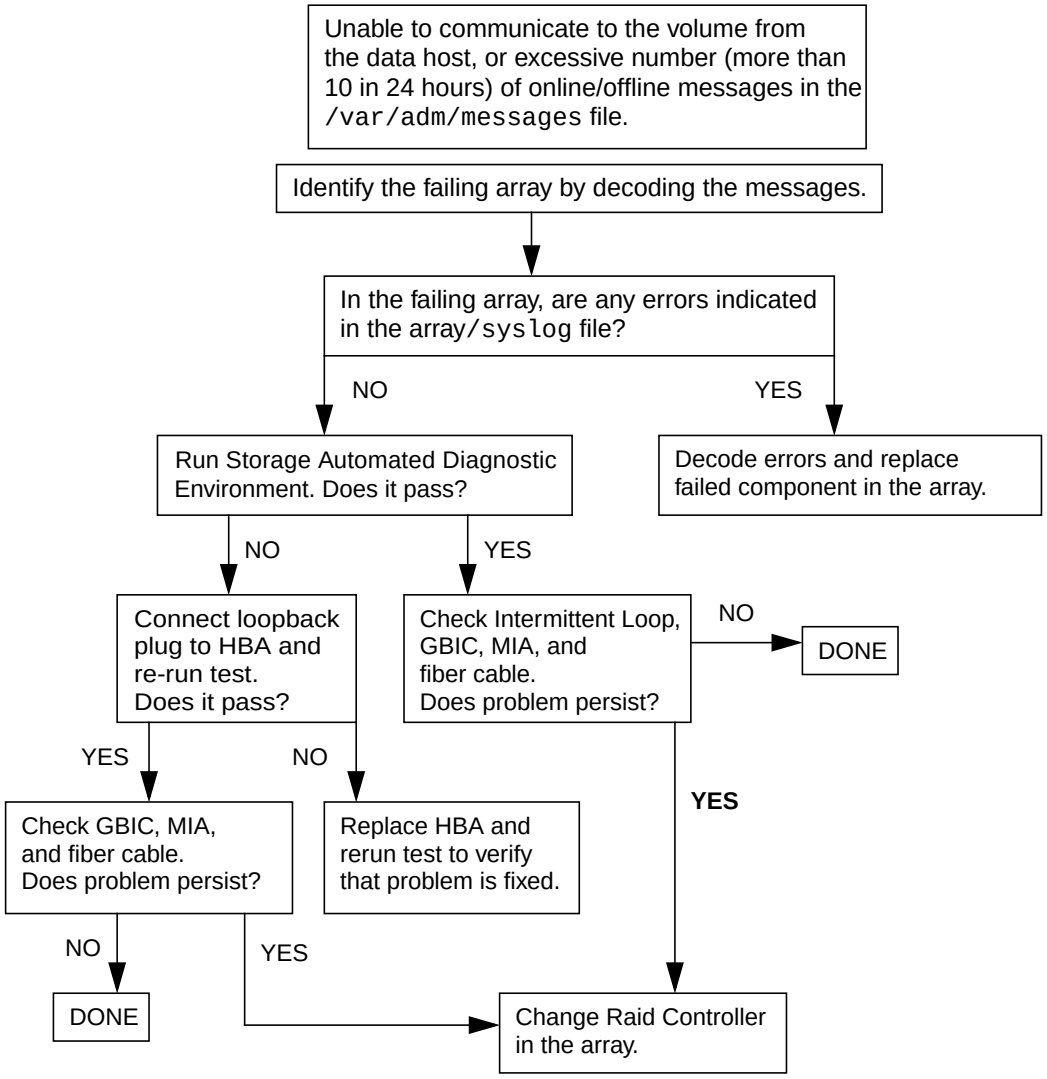

FIGURE 3-1 Data Connection Troubleshooting Flow Chart

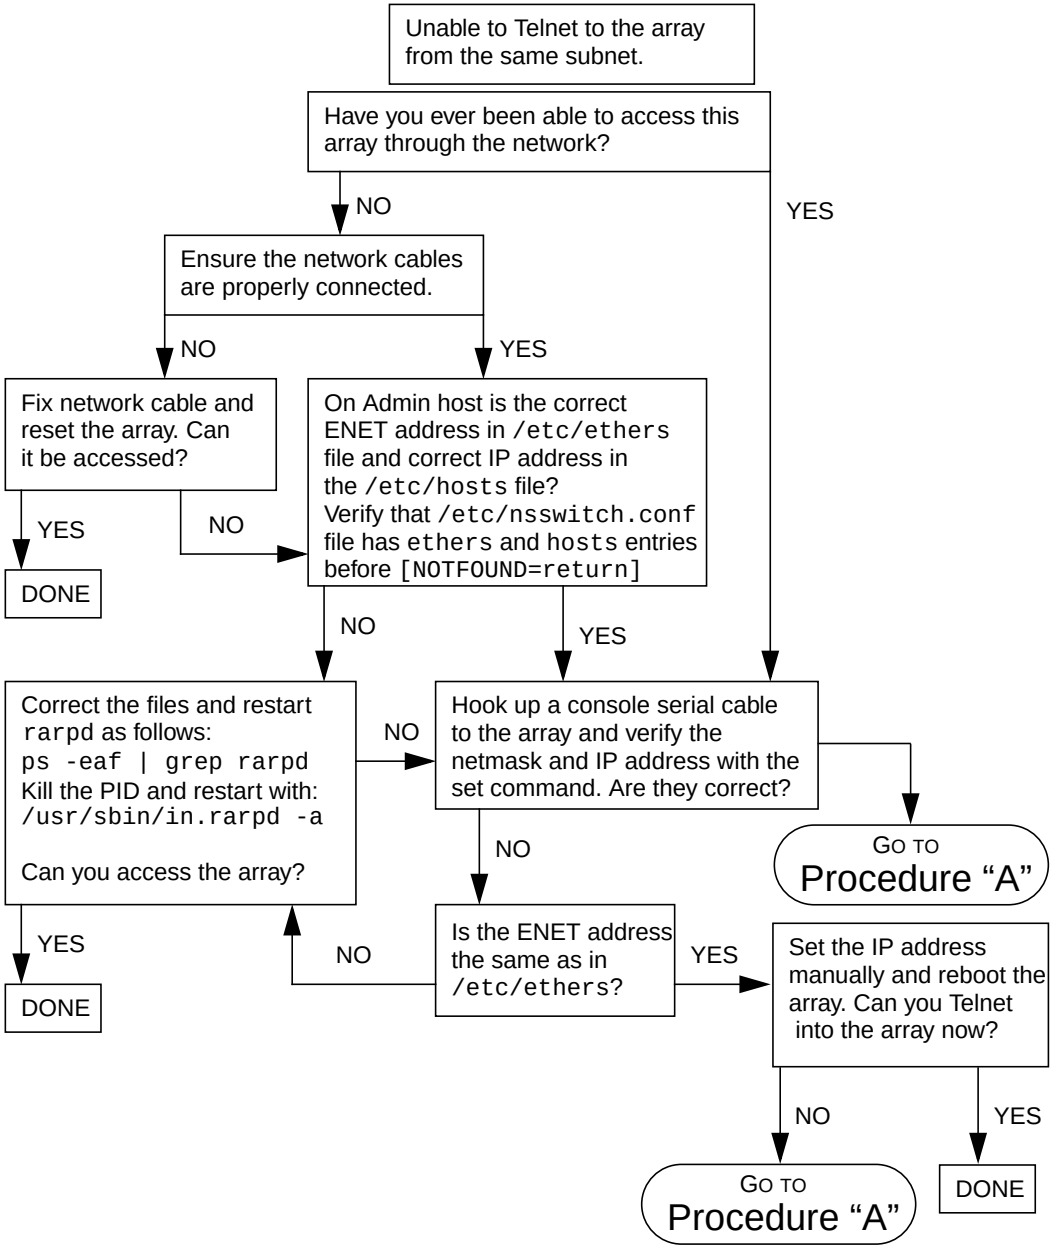

FIGURE 3-2 Ethernet Troubleshooting Flow Chart

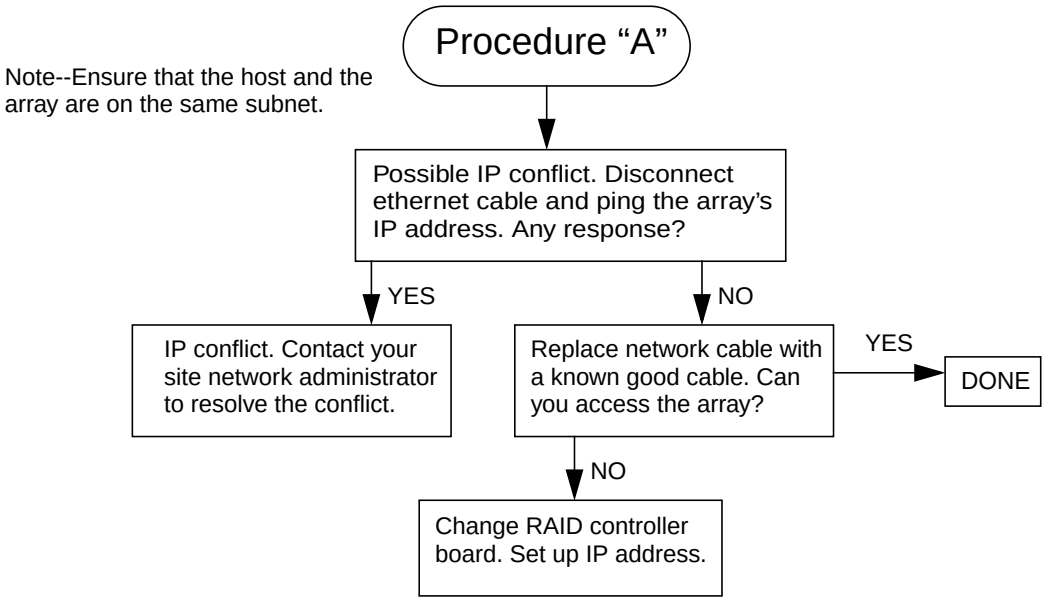

FIGURE 3-3 Procedure “A”

Initial Troubleshooting Guidelines

To begin a problem analysis, check one or more of the following information sources for troubleshooting and/or perform one or more of the following checks.

Troubleshooting Sources

1. **The array LEDs, which provide a visual status as described in** *Sun StorEdge T3+ Array Installation and Configuration Manual***.**
2. **Sun StorEdge T3+ array generated messages, found in a log file, indicating a problem or system status with the array. See "Sun StorEdge T3+ Array Generated Messages" on page 2 for more information about array generated messages.**
3. **Host-generated message, found in the** `/var/adm/messages` **file, indicating a problem with the host channel connection to the array unit. See "Host-Generated Message" on page 2 for more information about host generated messages.**

Troubleshooting Checks

1. **The connection between the host and the Sun StorEdge T3+ array as described in "Storage Automated Diagnostic Environment Link Test" on page 27.**
2. **The array boot status as described in "Checking Array Boot Status" on page 27.**
3. **FRU status as described in the** *Sun StorEdge T3+ Array Administrator's Manual***.**
4. **Array status as described in "Telnet Connection Status Checks" on page 30.**
5. **Array operation as described in "Testing the Array With Storage Automated Diagnostic Environment" on page 36.**
6. **Miscabled partner groups as described in "Identifying Miscabled Partner Groups" on page 36.**
7. **Data channel as described in "Identifying Data Channel Failures" on page 39.**

Verifying the Data Host Connection

To verify the physical connection between the host and the array, use a utility such as the `format` command in the Solaris environment. The output of the command confirms whether a volume is on the array. For example:

- **On the application host, enter** `format` **at the supervisor prompt.**

```
# format
Searching for disks...done

c1t1d0: configured with capacity of 133.38GB

AVAILABLE DISK SELECTIONS:
       0. c0t2d0 <drive type unknown>
          /sbus@1f,0/SUNW,fas@e,8800000/sd@2,0
       1. c0t3d0 <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
          /sbus@1f,0/SUNW,fas@e,8800000/sd@3,0
       2. c1t1d0 <SUN-T3-0100 cyl 34145 alt 2 hd 64 sec 128>
          /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w50020f2300000121,0
Specify disk (enter its number):
```

In this example, device number 2 is a volume on the array, as identified by the `SUN-T3-0100` label.

Storage Automated Diagnostic Environment Link Test

Use the Storage Automated Diagnostic Environment to verify the physical connection between the host, array, and any other physical devices and to determine the primary and alternate paths. Access the Storage Automated Diagnostic Environment main window and click the Diagnose link. Then click the Diagnostics Tests link. See the *Storage Automated Diagnostic Environment User's Guide* for instructions.

Caution – Any Sun StorEdge T3+ Array that is connected to a host via a switch by using F-Ports on the array side will fail. If the port is an F-Port, you need to remove the cable from the array before running Switchtest. The link, however, works if the array is configured on a TL-Port. This note is in reference to BugID 4731718.

Checking Array Boot Status

Establish a serial connection with the array as described in "Establishing a Serial Port Connection" on page 7.

8. **Boot the array.**

   There are several ways to initiate a boot cycle:

   - Power off the Sun StorEdge T3+ array and power it on again.
   - Log into a array and issue a `reset` command.
   - Log into a array and issue a `shutdown` command (this requires a power cycle to get the system to start booting).
   - If the array is already running, you can reboot by issuing a `boot` command with options.

Screen messages similar to the following appear:

```
:/> tip -9600 /dev/ttyb

Connected.

T3-EP Release 2.01 2002/3/8 13:05:27 (IP Address of tray)
Copyright (c) 1997-1999 Sun Microsystems, Inc.
All Rights Reserved
Found units:   u1-ctr1  [u2-ctr]
auto boot is enabled
hit the RETURN key within 3 seconds to cancel...
Default master is 1
Default alternate master is 2
Initializing System Drives...
Initializing XPT Components...
Initializing QLCF Components
Initializing Loop 1 ISP2100...firmware status 3
Detecting 20 FC-AL ports on loop 1
Initializing Loop 2 ISP2100...firmware status 3
Detecting 20 FC-AL ports on loop 2
Initializing SVD Services...
Found (18) disks in the system
Found 9 disks in U1...
Found 9 disks in U2...
Trying to boot from encid 1...
Booting from U1D1...
Executing...
Starting...

Login: root
password: root-password
```

Once the array starts a full boot, any system problems detected by the system are displayed. pSOSystem also identifies any controllers not responding, or if the master has failed over to the alternate master, in the boot messages.

Firmware status codes are good indicators of internally detected system and configuration problems. In the previous `boot` message example, a firmware status of 3 is displayed. This status implies the array is ready for operation. TABLE 3-2 lists other firmware status codes that can be reported through the serial-port console during a array `boot` cycle.

TABLE 3-2 Firmware Status Indicators

Status	Definition
0	ISP is waiting for configuration process to complete
1	ISP is waiting for ALPA assignment
2	ISP is waiting for port login
3	ISP is ready and optimal
4	ISP has lost loop synchronization
5	ISP has experienced an unrecoverable error
6	Reserved
7	ISP is not participating on the loop

Once the array has fully booted, all the commands available through the CLI are accessible.

Note – If you make configuration changes at the EPROM prompt, they can be overwritten when the array boots completely. Check the array settings after the array has booted to ensure that they are correct.

A message such as the following might appear after you log in:

```
6.1/: device not mounted
```

It is possible that the serial cable is connected to the alternate master unit instead of the master unit. To determine if this is true, enter `fru myuid` and SYS STAT.

```
:/: fru myuid

u1
```

If you are connected to the alternate, stop the tip session, reconnect the serial cable to the master unit, and start the session again. Verify that the role of the unit to which you are connected is specified as “Master.”

Telnet Connection Status Checks

Check array status using a variety of CLI commands. This section contains the following topics:

Determining Failover

1. **On the host, use the** `telnet` **command with the array name (or IP address) to connect to the array.**

```
mngt_host# telnet disk-tray-name
Trying 172.20.57.30...
Connected to auggie.Central.Sun.COM.
Escape character is '^]'.

Telnet session (172.20.57.30)
```

2. **Log in to the array by typing** `root` **and the supervisor password at the prompts.**

3. **To determine which unit is the master or alternate master unit, enter** `sys stat`**.**

   The following example shows a partner group in a normal state:

```
:/:<2>sys stat
Unit   State      Role    Partner
-----  ---------  ------  -------
 1     ONLINE     Master    2
 2     ONLINE     AlterM    1
```

   In a failover state, unit 2 assumes the role of master controller and unit 1 is disabled, as shown in the following example:

```
:/:<3>sys stat
Unit   State      Role    Partner
-----  ---------  ------  -------
  1    DISABLED   Slave
  2    ONLINE     Master
```

4. **Use the** `port list` **command to display how paths are mapped from the host ports to the volume.**

   This displays World Wide Names (WWNs) that can be compared to the WWNs displayed by the Solaris command `format`(1M).

```
:/:<4>port list

port   targetid   addr_type   status   host   wwn
u1p1       1         hard      online    sun    50020f23000002ba
u2p1       2         hard      online    sun    50020f23000002cd

mgmt-host# format
       Searching for disks...done

AVAILABLE DISK SELECTIONS:
0. c0t0d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
    /pci@1f,4000/scsi@3/sd@0,0
1. c2t1d0 <SUN-T300-0101 cyl 34145 alt 2 hd 64 sec 128>
    /pci@6,2000/SUNW,ifp@1/ssd@w50020f23000002ba,0
Specify disk (enter its number):
```

   In the example above, the WWN of 50020f23000002ba identifies the port and volume match.

Verifying the Firmware Level and Configuration

The Sun StorEdge T3+ array has four different types of firmware:

- Controller firmware. See the *Sun StorEdge T3+ Array Installation and Configuration Manual*.
- Interconnect card firmware. See the *Sun StorEdge T3+ Array Installation and Configuration Manual*.
- Controller electrically erasable programmable read-only memory (EPROM) firmware. See "Controller EPROM Firmware" on page 51".
- Disk drive firmware. See "Check the drive status to ensure that the reconstruction of the replaced drive FRU has completed." on page 70".

The firmware upgrade procedures must be done through the Ethernet connection. The latest firmware versions are located on the SunSolve℠ web site:

```
http://sunsolve.sun.com
```

The current firmware file naming restrictions are as follows:

- The name consists of a string of 1 to 12 characters.
- The name must start with an alphabetic character and not a numeral. For example:
  - `file1.bin` is acceptable
  - `1file.bin` is not acceptable
- The characters can be a combination of the following:
  - alphabetic letters
  - digits (0 through 9)
- Special characters such as:
  - _ (underscore)
  - . (period)
  - $ (dollar symbol)
  - - (dash).
- Names are case-sensitive. (For example, `ABC` and `abc` are different files.)

Make sure the latest firmware versions are installed and that the array configuration information indicates that the unit is ready for operation.

Check the firmware versions and array information in a telnet session with the array.

1. **On the host, use the** `telnet` **command with the array name (or IP address) to connect to the array.**

   For example:

```
# telnet array-name
Trying 23.123.123.3...
Connected to 123.123.123.3.
Escape character is '^]'.

Telnet session (123.123.123.3)
```

2. **Log in to the array by typing** `root` **and the supervisor password at the prompts.**

   The array prompt is displayed.

3. **Enter** `ver` **to identify the controller firmware.**

   For example:

```
:/:<5>ver

T3B Release 2.01.01 2002/07/30 19:16:42 (10.4.35.134)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.
```

   The `ver` command displays the header information.

4. **Enter `fru list` to display the firmware for the disk drives, interconnect card, and EPROM level.**

   In the event of a FRU failure, `fru list` output contains the serial numbers helpful in verifying correct FRU replacement.

```
:/:<7>fru list
ID      TYPE                VENDOR       MODEL          REVISION       SERIAL
------  ----------------    -----------  -----------    -------------  --------
u1ctr   controller card     0301         501-5710-02(   020100/020101  112035
u2ctr   controller card     0301         501-5710-02(   020100/020101  112122
u1d1    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HMKJ
u1d2    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HH2A
u1d3    disk drive          SEAGATE      ST336704FSUN   A726           3CD1H9WS
u1d4    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HM64
u1d5    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HMC2
u1d6    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HM63
u1d7    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HE3A
u1d8    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HNK0
u1d9    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HM5P
u2d1    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HHH5
u2d2    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HMJC
u2d3    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HGKR
u2d4    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HLBJ
u2d5    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HNH0
u2d6    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HH4Z
u2d7    disk drive          SEAGATE      ST336704FSUN   A726           3CD1H92W
u2d8    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HN9T
u2d9    disk drive          SEAGATE      ST336704FSUN   A726           3CD1HKQP
u1l1    loop card           SCI-SJ       375-0085-01-   5.02 Flash     1413
u1l2    loop card           SCI-SJ       375-0085-01-   5.02 Flash     2294
u2l1    loop card           SCI-SJ       375-0085-01-   5.02 Flash     001415
u2l2    loop card           SCI-SJ       375-0085-01-   5.02 Flash     002054
u1pcu1  power/cooling unit  TECTROL-CAN   300-1454-01(   0000           001787
u1pcu2  power/cooling unit  TECTROL-CAN   300-1454-01(   0000           001784
u2pcu1  power/cooling unit  TECTROL-CAN   300-1454-01(   0000           001544
u2pcu2  power/cooling unit  TECTROL-CAN   300-1454-01(   0000           001545
u1mpn   mid plane           SCI-SJ       370-3990-01-   0000           000953
u2mpn   mid plane           SCI-SJ       370-3990-01-   0000           000958
```

In this example:

- EPROM firmware version is `Controller card, Revision 020100/020101`
- Disk drive firmware version is `Revision A726`
- Interconnect card (`loop card`) firmware version is `Revision 5.02 Flash`

Checking FRU Status

Use the `fru stat` command to provide a status of each FRU, including temperatures.

```
:/:<43>fru stat
CTLR    STATUS   STATE       ROLE        PARTNER    TEMP
------  -------  ----------  ----------  -------    ----
u1ctr   ready    enabled     master      u2ctr      31.0
u2ctr   ready    enabled     alt master  u1ctr      30.5

DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      30    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      29    vol1
u1d6    ready    enabled     data disk   ready      ready      30    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      32    vol3
u2d1    ready    enabled     data disk   ready      ready      34    vol2
u2d2    ready    enabled     data disk   ready      ready      38    vol2
u2d3    ready    enabled     data disk   ready      ready      36    vol2
u2d4    ready    enabled     data disk   ready      ready      37    vol2
u2d5    ready    enabled     data disk   ready      ready      34    vol2
u2d6    ready    enabled     data disk   ready      ready      36    vol4
u2d7    ready    enabled     data disk   ready      ready      35    vol4
u2d8    ready    enabled     data disk   ready      ready      40    vol4
u2d9    ready    enabled     data disk   ready      ready      36    vol4

LOOP    STATUS   STATE       MODE        CABLE1     CABLE2     TEMP
------  -------  ----------  -------     ---------  ---------  ----
u2l1    ready    enabled     master      installed  -          29.5
u2l2    ready    enabled     slave       installed  -          31.0
u1l1    ready    enabled     master      -          installed  29.5
u1l2    ready    enabled     slave       -          installed  30.5

POWER   STATUS   STATE       SOURCE  OUTPUT  BATTERY  TEMP    FAN1    FAN2
------  -------  ---------   ------  ------  -------  ------  ------  ------
u1pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u1pcu2  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu2  ready    enabled     line    normal  normal   normal  normal  normal
```

Note – The `fru stat` command reports temperature readings on the interconnect cards, controller board, disk drives, and PCUs. For the PCU, the `fru stat` output does not display a numeric temperature, but instead reports a temperature state. For all other FRUs, `fru stat` reports a numerical temperature.
System firmware monitors only the temperature state reported by the PCUs. This means a high temperature reading on an interconnect card, for example, will not cause the firmware to take evasive action (such as powering off the array).

Testing the Array With Storage Automated Diagnostic Environment

Access the Storage Automated Diagnostic Environment main window and click the Diagnose link. Then click the Diagnostics Tests link. See the *Storage Automated Diagnostic Environment User's Guide* for instructions.

Identifying Miscabled Partner Groups

If a partner group has booted successfully but is unable to establish a Telnet connection with the management host, the partner group might be cabled together incorrectly.

The interconnect cable connections between dual controller units are critical for determining which unit is the master controller and which is the alternate master. If the interconnect cables are not properly installed on the interconnect cards, the top unit could boot as the master controller, and the bottom unit would assume alternate master status. Because the host has been configured to use the MAC address of the bottom unit, this alternate configuration can cause the units to malfunction.

If the bottom unit is incorrectly cabled, making the bottom unit the alternate master, the bottom unit's Ethernet port will be inactive unless a fail over situation occurs. In that event, the IP and MAC address of the bottom unit will take over the values of the master (top) unit.

If the partner group has been cabled together incorrectly, the following procedure can help determine if the top unit is acting as the master controller.

1. **Determine the MAC address of the top unit.**

   The MAC address is located on a pull-out tab at the front of the unit, to the left of the first disk drive. (FIGURE 3-4).

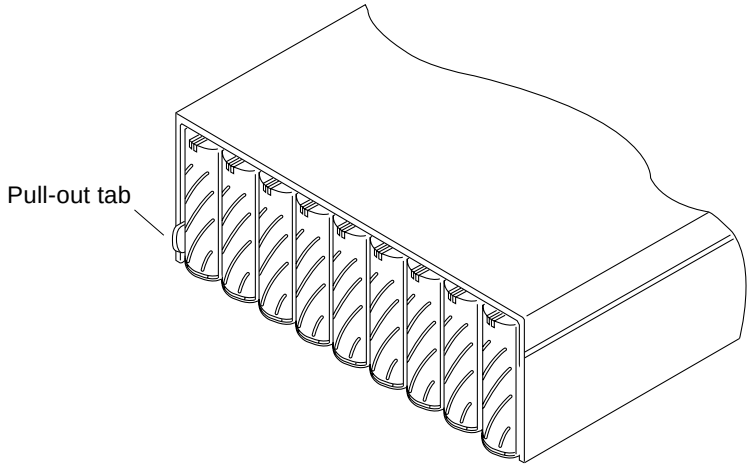

FIGURE 3-4 MAC Address on the Pull-Out Tab

2. **Edit the files on the RARP server to include the MAC address of the top unit.**

   a. **Edit the `/etc/ethers` file by adding the MAC address and array name.**

      For example:

```
8:0:20:7d:93:7e array name
```

      In this example:

      - `8:0:20:7d:93:7e` is the MAC address
      - *array name* is the name of the master controller unit.

   b. **Edit the `/etc/hosts` file with the IP address and array name.**

      For example:

```
123.123.123.111 array name
```

      In this example, `123.123.123.111` is the assigned IP address.

c. Edit the `/etc/nsswitch.conf` file to reference the local system files.

To ensure the Solaris software environment uses the changes made to `/etc/ethers` and `/etc/hosts` files, edit the `host` and `ethers` entries in the `/etc/nsswitch.conf` file so that the files parameter appears before the `[NOTFOUND=return]` statements as shown:

```
hosts:      nis files [NOTFOUND=return]
ethers:     nis files [NOTFOUND=return]
```

d. Determine if the RARP daemon is running by typing:

```
# ps -eaf | grep rarpd
```

- If the RARP daemon is running, proceed to Step 3.
- If the RARP daemon is not running, continue to Step e.

e. Start the RARP daemon in the Solaris environment by typing:

```
# /usr/sbin/in.rarpd -a &
```

3. **Ensure that there is an Ethernet connection to the 100BASE-T port of the top unit.**

4. **Press the power switch on the power and cooling units on both arrays to remove AC power (**FIGURE 3-5**).**

   It may take some time for the units to power off while shutdown procedures are performed. Wait until the units have powered off completely.

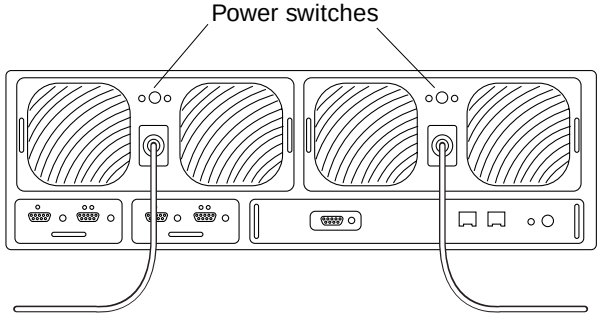

FIGURE 3-5 Power Switch Locations

5. **After both units have powered off, press the power switch on the power and cooling units again to restore power to and reset the arrays.**

   It may take up to several minutes for the arrays to power on and come back online. All LEDs will be green when the unit is fully powered on.

6. **After the units are fully powered on, start a Telnet session.**

   The Telnet session will connect to the top unit.

   If the host cannot `telnet` to the array, investigate the following other possible causes:

   - **RARP server not responding.** To determine if this is the problem:
     - Verify that the RARP daemon is running on the host system.
     - Verify that the `/etc/nsswitch.conf` file is properly configured on the RARP server.
     - In the Solaris environment, use the `snoop` command to verify that the array is attempting to establish RARP communication with the Solaris server.
   - **MAC address is incorrect.** In the Solaris environment, use the `snoop` command to specify the MAC address of the array and to determine if any RARP packets are transmitted. If you observe no transmissions during a reboot of the array, verify that the MAC address on the array label matches the MAC address configured on the RARP server.
   - **Netmask is incorrect.** The default netmask address used on the array is 255.255.255.0. If the local subnet uses a different netmask, the RARP operation might not work.
   - **Inoperable network connections.** If using hubs to connect to the network, try eliminating or replacing the hub.
   - **Incorrect IP address.** Connect to the array through the serial port, and verify that the IP address is correct.

Identifying Data Channel Failures

The data channel encompasses the host data path that extends from the host bus adapter to the media interface adapter (MIA) attached to the array. Errors in the host data channel are outside of the scope of the Sun StorEdge T3+ array. To determine failures in the data path, use the Storage Automated Diagnostic Environment.

Refer to the documentation of the selected diagnostics tool for information on identifying data channel failures.

Reserved System Area Recovery Procedure

Some of the conditions that indicate a corrupted system area of a Sun StorEdge T3+ array are:

- The controller is disabled or the booting process is cycling
- The command line prompt cannot be accessed using either the Ethernet or a serial interface.
- The application host cannot communicate with the LUNs.

Note – After configuring a system, always record the following data to prepare for the possibility of having to perform a recovery procedure:
- Array block size
- Multipathing settings
- Volume configuration
- Volume slicing configuration
- LUN masking settings

Recovery Procedure

1. **Establish a serial port connection to each Sun StorEdge T3+ array. See "Establishing a Serial Port Connection" on page 7.**
2. **Stop the application and unmount the file systems on the application host for the LUNs defined on the array(s) that are being recovered.**

```
# unmount filesystem
```

3. **Power off the affected array(s) by pushing both power switches on the PCUs.**
4. **Clear the controller disable flags by partially removing all Interconnect cards half way out for 30 seconds. Then re-install the Interconnect cards.**
5. **If configured as an enterprise configuration, disconnect the interconnect cables from the alternate master controller. See** FIGURE 3-6**.**

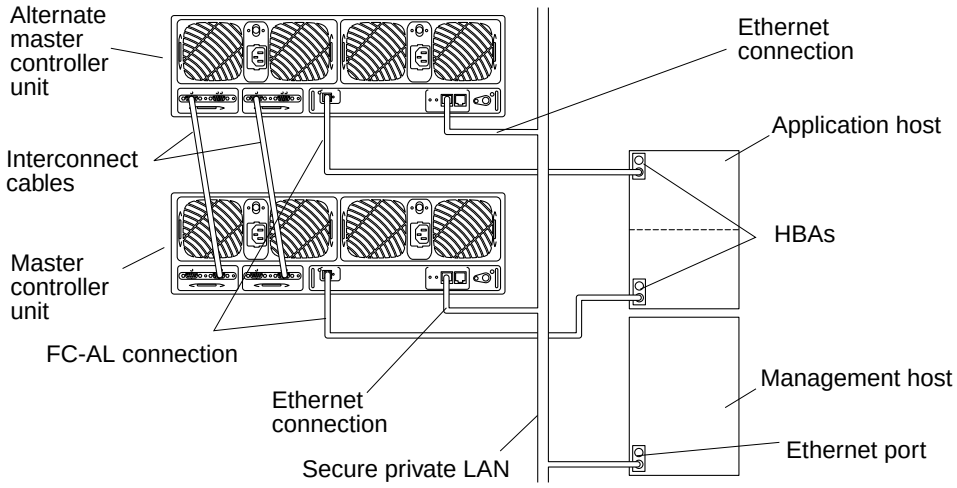

FIGURE 3-6 Single Host With Two Controller Units Configured as a Partner Group

6. **Power on the array or the master controller unit of an enterprise configuration.**

   The array starts to boot automatically.

7. **Stop the boot process at the cancellation message by pressing the Return key.**

```
T3B-EP Release 2.00 2001/06/22 16:07:00 (172.20.57.31)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.

Found units: [u1-ctr] u2-ctr
tftp boot is enabled
hit the RETURN key within 3 seconds to cancel...Cancelled!
T3B-EP>
```

8. **Set the array to boot from the tftp boot server. See "Using `tftpboot` to Boot a Single Array or a Partner Group Remotely" on page 13.**

   Verify that the bootmode, tftp host and tftp file settings are correct.

```
T300-EP> set bootmode tftp
T300-EP> set tftphost IP_address
T300-EP> Set tftpfile controller_binary
T300-EP> set
bootmode  tftp
bootdelay 3
sn        xxxxxx
ip        10.1.102.112
netmask   255.255.255.0
gateway   xxx.xxx.xxx.xxx
tftphost  xxx.xxx.xxx.xxx
tftpfile  nb210.bin
hostname  T3
spindelay 0
revision  0210
mac       xx:xx:xx:xx:xx
rarp      on
```

9. **Reset the array.**

```
T300-EP> reset
```

10. **Observe the boot cycle.**
    - If the system is able to boot to a normal login prompt, proceed to Step 11.
    - If the array continues to boot in a cycle, stop the cycle and break to the diagnostic menu by pressing Ctrl-t and continue pressing at one second intervals until the booting stops. Press Return at the diagnostic menu prompt and continue below:

      i. From the diagnostic menu select "Quit but go into Label Control Menu."

      ii. From the Label Control Menu, select "Wipe out unit 1 Sysarea LFS."

      iii. Select "Quit All."

      The system should continue the boot cycle.

      iv. Verify the system boots to the login prompt and log in as the root user.

v. Use the appropriate patch to execute the `t3.sh` script to restore the missing files to the array local file system.

- Sun StorEdge T3 controller - patch 109115
- Sun StorEdge T3+ controller - patch 112276

11. Install the boot code by typing:

```
T3:/:<1> boot -i nb210.bin
```

12. Set the boot mode to auto by typing:

```
T3:/:<2> set bootmode auto
```

13. Reset the array by typing:

```
T3:/:<3> reset -y
```

Verify the system boots normally by observing the console.

- If you have a workgroup configuration, proceed to Step 14.
- If you have an enterprise configuration, continue below:

i. Shutdown the array by typing:

```
T3:/:<4> shutdown -y
```

ii. Power off the array.

iii. Re-attach the interconnect cables.

iv. Power on both arrays of the enterprise configuration.

v. Verify the systems boot normally by observing the console of each controller.

14. **Verify that the system parameters are set correctly by typing:**

```
T3:/:<5> sys list
blocksize          : 16k
cache              : off
mirror             : off
mp_support         : none
naca               : off
rd_ahead           : on
recon_rate         : med
sys memsize        : 128 MBytes
cache memsize      : 1024 MBytes
enable_volslice    : on
fc_topology        : auto
```

Caution –

- Failure to ensure that the blocksize is set correctly will lead to data loss or corruption.
- Failure to ensure multipathing support is enabled will prevent proper LUN failover in an enterprise configuration.
- Failure to restore volume slices on the correct blocks will cause data loss or corruption.
- Failure to ensure LUN masking is properly restored can result in data inaccessibility on the desired host or result in improper access from undesired hosts.

15. **If the volume information was lost, add the array volumes using the same geometry in which they were previously created.**

    Ignore any SVD_PATH_FAILOVER or SVD_CHECK_ERROR messages that occur.

```
T3:/:<6> vol add vol1 data u1d1-8 raid 5 standby u1d9
T3:/:<8> vol add vol2 data u2d1-8 raid 5 standby u2d9
T3:/:<9> vol stat
vol1          u1d1   u1d2   u1d3   u1d4   u1d5   u1d6   u1d7   u1d8   u1d9
unmounted      0      0      0      0      0      0      0      0      0
vol2          u2d1   u2d2   u2d3   u2d4   u2d5   u2d6   u2d7   u2d8   u2d9
unmounted      0      0      0      0      0      0      0      0      0
```

16. **Fast initialize the array volumes by typing:**

```
T3:/:<10> .vol init vol1 fast
WARNING - Existing volume data won't be changed.
Continue ? [N]: y
T3:/:<11> .vol init vol2 fast
WARNING - Existing volume data won't be changed.
Continue ? [N]: y
```

17. **Mount the array volumes.**

```
T3:/:<12> vol mount vol1
T3:/:<13> vol mount vol2
T3:/:<14> vol stat
vol1          u1d1   u1d2   u1d3   u1d4   u1d5   u1d6   u1d7   u1d8   u1d9
mounted        0      0      0      0      0      0      0      0      0
vol2          u2d1   u2d2   u2d3   u2d4   u2d5   u2d6   u2d7   u2d8   u2d9
mounted        0      0      0      0      0      0      0      0      0
```

18. **Enable volume slicing if applicable and restore the slices as they previously existed.**

19. **Restore the LUN masking settings on the volume slices as applicable.**

20. **Verify that the application host can access the restored array LUNs by typing:**

```
# luxadm probe
```

21. **Rescan the devices with Volume Manager, if applicable by typing:**

```
# vxdctl enable
```

22. **Check the file systems on the appropriate LUNs by typing:**

```
# fsck filesystem
```

23. **Mount the file systems and restart the application by typing:**

```
# mount filesystem
```

24. **Create a `syslog.conf` file with the correct remote and local logging entries. Upload it to the array by using ftp and place it in the `/etc` directory.**

25. **Restart the Sun StorEdge T3+ array syslog daemon.**

```
T3:/:<15> set logto *
```

26. **Use the `logger` command to verify the system is logging properly by typing:**

```
T3:/:<16> logger message
```

    where *message* is the text of a test message to be logged.

27. **Create a schd.conf file with the correct BAT_BEG date and 28-day BAT_CYC. Refer to "Note that the next refresh start time is always calculated from the start time of the previous refresh cycle. If a user manually starts a refresh cycle, then the next refresh depends on the starting time of the manually activated refresh cycle." on page 89.**

    Make sure to specify a future start date to preclude a refresh during the recovery.

28. **Restart the battery scheduler.**

29. **Verify the battery scheduler is working as expected by typing:**

```
T3:/:<16> refresh -i
T3:/:<17> refresh -s
```

30. **Exit from the serial console session on each controller.**

CHAPTER 4

Controller Card Assembly

This chapter describes how to monitor and replace the controller card and how to upgrade the firmware. The chapter contains the following sections:

Controller Card LEDs

This section describes the controller card LEDs for the Sun StorEdge T3+ array.

Sun StorEdge T3+ Array Controller Card LEDs

The Sun StorEdge T3+ array controller card has two channel-active LEDs, one for the FC-AL interface port and one for the Ethernet port, and a controller status (online) LED. TABLE 4-1 lists the possible states of the controller card LEDs and describes each state.

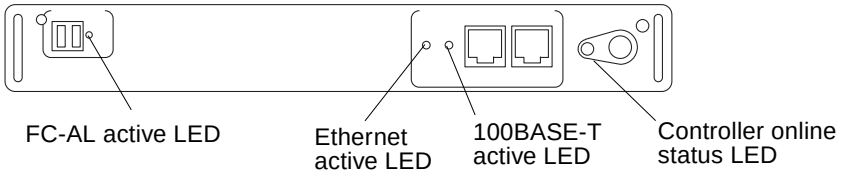

FIGURE 4-1 Sun StorEdge T3+ Array Controller Card LEDs

TABLE 4-1 Sun StorEdge T3+ Array Controller Card LED Descriptions

LED	Action	Description
FC-AL Channel Active LED (green)	Off	Port disabled
	Green	Port enabled and idle
	Blinking green	Port enabled and active
Ethernet Active LED (green)	Off	Link invalid
	Green	Link valid and idle
	Blinking green	Link valid and active
100 BASE-T Active LED (green)	Off	Port disabled (10 Mbps rate)
	Green	Port enabled and idle (100 Mbps rate)
	Blinking green	Port enabled and active
Controller Status LED (green or amber)	Off	Controller not installed (not recognized)
	Green	Controller OK
	Amber	Controller boot, shutdown, or firmware download in progress
	Blinking amber	Controller failure; OK to replace controller

Note—Verify a controller card failure using the CLI.

Removing and Replacing a Controller Card

Note – A new feature of the version 2.0 controller firmware is Autoversioning. This feature allows you to seamlessly update from a Sun StorEdge T3 array to a Sun StorEdge t3+ array. When a Controller card is replaced, Autoversioning ensures that the new controller is flashed with the latest firmware version of the existing array controller of an enterprise configuration and that both controllers are therefore running the same firmware version.

A controller card can be replaced without system interruption *only* if the array is configured in a partner group (redundant controller unit configuration).

Caution – A removed controller card must be replaced within 30 minutes or the Sun StorEdge T3+ array and all attached arrays will automatically shut down and power off.

To replace the controller card:

1. **Observe static electricity precautions.**

   See “Static Electricity Precautions” on page 5.

2. **Ensure that the controller card is showing failure status.**
3. **Remove the Ethernet cable from the 100BASE-T connector.**
4. **Remove the fiber-optic cable (and MIA, if applicable) from the FC-AL connector.**
5. **Unlock the controller card by pushing in on the latch handles.**

   Use a coin or small screwdriver to press in and release the latch handle.

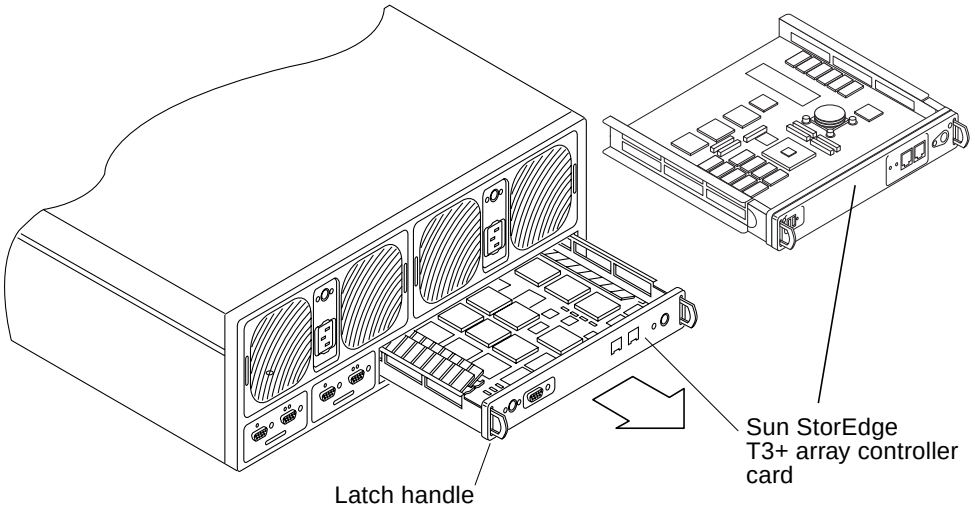

FIGURE 4-2 Removing the Controller Card

6. **Pull the controller card out using the latch handles.**
7. **Insert the new controller card.**
8. **Lock the new controller card by pushing in the latch handles.**

   Use a coin or small screwdriver to press in and secure the latch handle
9. **Insert the fiber-optic cable (and MIA for T3 controllers) back into the FC-AL connector.**
10. **Insert the Ethernet cable into the Ethernet port.**
11. **Check the controller status LED to determine when the controller is operational.**

    While the controller boots, the controller status LED is solid amber. When the controller is operational, the LED is green.
12. **Verify the status of the controller card using the CLI.**

    Refer to “Checking FRU Status” on page 35 for instructions.

Note – In a partner group configuration, the controller fails over to the alternate master when there is a controller card failure in a master unit. After the controller board is replaced, use the `reset` command if you wish to have u1 become the master again.

Upgrading Controller Firmware

The controller firmware can be upgraded on an operational system. However, for the upgrade to take effect, the controller must be reset (booted). While the controller boots, the array is not available for storage.

The firmware upgrade procedures that follow must be done through the Ethernet connection. The latest firmware version is located on the SunSolve web site:

```
http://sunsolve.sun.com
```

The following conditions apply to firmware upgrades:

- The firmware has to be resident on the host for this operation.
- The Sun StorEdge T3+ array has to have a *root password* prior to attempting this procedure.

To upgrade the firmware, see the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

Controller EPROM Firmware

The EPROM firmware is stored in the FLASH memory device (FMD) on the controller card. The array can be operational during the EPROM firmware upgrade.

Note – To upgrade the EPROM firmware in a partner group, you need to perform this procedure only once for both units to be upgraded.

The latest firmware versions are located on the SunSolve web site:

```
http://sunsolve.sun.com
```

Firmware is released as a patch which consists of an entire `tar` file with an automated uploader script that copies the files (including the `ep` and `lpc` images) to the Sun StorEdge T3+ array being upgraded.

Firmware Upgrade Discussion

Boot Code Explanation

There are three levels of boot code plus an extended post code for factory testing.

- The first level selects and jumps to one of the two copies of the second level code. The second level code initializes memory and loads itself to RAM locations starting from 0x500000.
- The second level code can allow tftp boot or ROM boot for the third level code. In ROM boot, the second level code selects one of the two copies of the third level (RAID application) code. The second level code loads the RAID application code to RAM locations from 0x20000.
- The third level is the RAID application.
- The extended Power On Self Tests (post code) is for performing factory level diagnostics.

First Level Boot Code

The level 1 boot code starts at 0xFFF00100 which is the processor's reset vector. The first level code initializes the MPC107 bridge chip and the console serial port. It prints `T3B-` when the initialization is done. Then it waits about 1.5 seconds to allow the user to select one of the two copies of level 2 code to boot. The user can type `1` or `2`, but there is no echo for the character typed. If the user makes no selection, level 2 code does the selection automatically.

In the automatic selection, the level 1 code verifies the level 2 boot code stored in ROM. It finds which one is newer and jumps to the selected code.

If the user has entered the selection before automatic selection, level 1 code jumps to the one the user has selected after verifying the code is valid. If the user selects an invalid copy, then level 1 code jumps to the valid one instead of the user selected one.

After the level 1 code has decided which copy of level 2 code to use, it prints `1` or `2` to notify the user which copy is selected as a feedback.

The level 1 code starts at 0xFFF00100 and extends to 0xFFF20000.

If there is no valid level 2 code, the code prints `0` after `T3B-` and reboot.

Second Level Boot Code

The second level boot code is comparable to the EPROM mode of the Sun StorEdge T3 Array except the auto bootmode will boot from ROM instead of from disks. The level 2 code allows the user to:

- Set basic system configuration. For example, IP, Gateway, and bootmode.
- Allow `tftpboot` for the level 3 code.
- Allow `autoboot` for the level 3 code.
- Allow update of the ROM code when bootmode is `tftp`.

Note – Use the login prompt to set the IP address, netmask, and hostname instead of using the EP prompt. Setting these parameters at the EPROM level will be lost.

There are two copies of level 2 code; one in 0xFFF40000-0xFFF9FFFF, the other in 0xFFFA0000-0xFFFFFFFF.

To update the EPROM, the following commands should be used.

- `.ep download` *filename*

  This updates the level 1 boot code. Since there is only one copy of level 1 code in ROM, if this update fails, the controller may not be able to boot after the failure.

- `.ep download` *filename*

  This updates the level 2 boot code. Two copies of level 2 boot code are kept in ROM. Level 1 boot code selects the most recent updated one to boot. The user can override the automatic selection by typing `1` or `2` within 1.5 seconds after `T3B-` is displayed on the console.

  If the command fails in the middle of update, it will be an invalid level 2 code, and level 1 code will not select the invalid level 2 code for booting.

  If a bad level 2 code is programmed into ROM successfully, then the user can manually select which copy to boot in order to work around the bad level 2 code. If this happens, it is better to update the level 2 code again in order to override the bad level 2 code copy.

  Level 2 code has a size limitation of 384 Kbytes.

During boot up, the level 2 code occupies RAM space starting at 0x500000, and the level 3 code is loaded by the level 2 code. Currently the starting location of level 3 code is fixed at 0x20000. Although level 3 code can start at another location, the code space after upload to RAM cannot go over 0x500000.

The network of level 2 code will be enabled only when the bootmode is set to “tftp”. Thus, the `.ep` command will only work when bootmode is “tftp”.

The level 2 code also includes POST (Power-On Self Test) code in the booting process.

Third Level Boot Code

The third level boot code is the RAID application. The code has assumed that level 2 code would have set up the MPC 107 and cleared the RAM if it is cold-boot.

There are two copies of level 3 code in ROM; one in 0xFF800000-0xFFB7FFFF, the other in 0xFFB80000- 0xFFF00000.

EPROM and tftp Download File

The file to be downloaded into ROM or through tftp must have specific header information with a structure such as listed below.

```
typedef struct ep_header_struct {
    init code_size;/* codesize */
    init code_cksum;/* codechecksum */
    init code_start;/* codestart */
    init code_base;/* codebase */

    init code_signature;/* codesignature */
    init code_rev;/* coderevision */
    init code_subrev;/* codesubrevision */
    init code_date;/* codedate */

    init code_time/* codetime */
    init hdr_counter;/* codecounter */
    init code_flags;/* codeflags */
    init reserved;

    init reserved[3];
    init hdr_cksum;/* headerchecksum */
    EP_HEADER;
```

The file content must be the binary image to be loaded into ROM or RAM. It cannot be an elf file, a hex file, or a srecord file.

The following explains each field in the header.

TABLE 4-2 Channel Active LED Descriptions

Header	Description
`code_size`	This is the size of the code without the header information. This value must be a multiple of four. The real file size should be `code_size` plus `sizeof` (EP_HEADER).
`code_cksum`	The 32-bit checksum value of the code. `code_cksum` = sum of all 32-bit words in code OR 0xFFFFFFFF+1.
`code_start`	The execution starting location. For example, after downloading level 3 code into RAM, the level 2 code jumps to this location.
`code_base`	The starting location for download. For example, the level 2 code downloads the level 3 code from ROM into RAM starting at this location. However, notice that when the starting location is 0x20000, the image in the file is loaded to 0x20040 in RAM. The first 0x40 bytes are occupied by the header information. That is to say that the `code_base` includes the space occupied by the header information.
`code_signature`	Each level of boot code has a unique signature. For example, the level 3 signature is “P2L3”.
`code_rev`	The revision of the code.
`code_subrev`	The subrevision of the code.
`code_date`	The date stamp of the code. For example, 20001225 means 2000/12/25.
`code_time`	The time stamp of the code. For example, 01020300 means 01:02:03.

TABLE 4-2 Channel Active LED Descriptions

Header	Description
`hdr_counter`	For the file to be downloaded to ROM by the `.ep` command to RAM through tftp, this field should be -1. But, after the code is programmed into ROM, the `.ep` command will change this field automatically. This field is used to identify which of two copies of level 2 code or level 3 code is newer. The smaller the value is, the older the code is. Thus, 0xFFFFFFFF is older than 0xFFFFFFFE. The `.ep` command will automatically update this field by taking the value of this field from the other copy and add 1 to the value.
`code_flags`	This field is used to identify whether special handling of the code file is needed. For example, the code file may be zipped, when it needs to be unzipped before uploading to RAM.
`hdr_cksum`	This is the checksum for the header portion. Since the `.ep` command will change `hdr_counter` when downloading code into ROM, this field will be updated accordingly by the `.ep` command.

Currently, only `code_signature` and `hdr_counter` affects the automatic selection done by level 1 code or level 2 code. Of course, `code_cksum` and `hdr_cksum` are used to validate the code.

Level 1 Controller Firmware

This procedure upgrades the level 1 firmware in only one controller. Therefore, you must perform it twice for each array enterprise configuration.

1. **Connect a console to the serial port (the right RJ-45 port) of the array.**

   See “Establishing a Serial Port Connection” on page 7.

Note – The serial cables used by the Sun StorEdge T3 and T3+ arrays are different. Both cables are supplied in the F370-4119-02 Diagnostic Kit. The T3 array uses the serial cable with RJ-11 connectors and the T3+ array uses the serial cable with RJ-45 connectors.

---

2. **Set up the tftp host.**

   See “Configuring a Server for Remote Booting” on page 16.

3. **Reset the controller and press any key on the console when the system prompt appears (within three seconds).**

4. **Set the bootmode and tftp settings as follows:**

```
T3B-EP> set bootmode tftp
T3B-EP> set tftphost xxx.xxx.xxx.xxx
```

Where *xxx.xxx.xxx.xxx* is the host IP address.

5. **Reset the T3+ system with the** `reset -y` **command.**

```
T3B-EP> reset -y
```

6. **Press a key from a serial port connection when the system prompts to press a key within three seconds.**

7. **Install the firmware using the** `.ep netload1` **command.**

```
T3B-EP> .ep netload1 level-1_image_filename
```

8. **Set the bootmode to automatic.**

```
T3B-EP> set bootmode auto
```

9. **Power cycle the array to reset it.**

   a. **Type:**

```
:/:<4> shutdown
shutdown the system, are you sure? [N] : y
```

   b. **Press the power button on each power and cooling unit to remove AC power.**

   c. **Press the power buttons again to return AC power to the array.**

Level 2 Controller Firmware

In an enterprise configuration, the `ep` command downloads level 2 firmware to both the master unit and alternate master unit at one time. To upgrade the Level 2 controller firmware perform the following steps.

1. **Use the ftp binary mode to transfer the firmware to the storage systems directory.**

   See “Establishing an FTP Session” on page 12.

2. **In a telnet session with the array, install the level 2 image. Type:**

```
:/:<1> .ep download level-2_image_filename
```

Level 3 Controller Firmware

In an enterprise configuration, this procedure downloads level 2 firmware to both the master unit and alternate master unit at one time. To upgrade the Level 3 controller firmware perform the following steps.

1. **Use the ftp binary mode to transfer the firmware to the storage systems directory.**

   See “Establishing an FTP Session” on page 12.

2. **In a telnet session with the array, set the** `bootmode` **to** `auto`**.**

```
:/:<2> set bootmode auto
```

3. **Install the level 3 image on the array.**

```
:/:<3> boot -i level-3_image_filename
```

4. **Reset the array.**

```
:/:<4> shutdown
shutdown the system, are you sure? [N] : y
```

   a. **Press the power button on each power and cooling unit to remove AC power.**

   b. **Press the power buttons again to return AC power to the array.**

Note – If during the boot process, a controller detects a level 3 firmware version on the system disk different than the level 3 image loaded in flash, the controller will reflash its local level 3 image and reset. This can appear as two sequential boot cycles. This process is expected behavior.

CHAPTER 5

Disks and Drives

This chapter describes how to monitor and replace the disk drives, upgrade the firmware, and repair corrupted disk labels. This chapter contains the following sections:

Monitoring Drive Status

The following sections describe commands for monitoring the status of the drives. Disk status can be checked by using a variety of CLI commands. This section discusses how to monitor the following:

1. **On the host, use the `telnet` command with the array name (or IP address) to connect to the array.**

```
mngt_host# telnet array name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

Telnet session (129.150.47.101)
```

2. **Log in to the array by typing `root` and the supervisor password at the prompts.**

Checking Drive Status Codes

- **Use the `vol stat` command to check drive status codes.**

  All drives should show a status of 0 under normal conditions.

```
:/:<40> vol stat

vol1          u1d1   u1d2   u1d3   u1d4   u1d5
mounted        0      0      0      0      0
vol2          u2d1   u2d2   u2d3   u2d4   u2d5
mounted        0      0      0      0      0
vol3          u1d6   u1d7   u1d8   u1d9
mounted        0      0      0      0
vol4          u2d6   u2d7   u2d8   u2d9
mounted        0      0      0      0
```

The following table lists numeric drive status codes.

TABLE 5-1 Drive Status Messages

Value	Description
0	Drive mounted
2	Drive present
3	Drive is spun up
4	Drive is disabled
5	Drive has been replaced
7	Invalid system area on drive
9	Drive not present
D	Drive is disabled and is possibly being reconstructed
S	Drive substituted (`vol recon` to stanby drive has completed)

Checking the Hot Spare

1. **Use the** `vol list` **command to check the location of the hot spare (standby) drive.**

```
:/:<41> vol list

volume          capacity     raid  data      standby
vol1            134.890 GB   5     u1d1-5    none
vol2            134.890 GB   5     u2d1-5    none
vol3            101.167 GB   5     u1d6-9    none
vol4            101.167 GB   5     u2d6-9    none
```

2. **Use the** `vol stat` **command to check the status of the hot spare drive.**

```
:/:<42> vol stat

vol1           u1d1   u1d2   u1d3   u1d4   u1d5
mounted         0      0      0      0      0
vol2           u2d1   u2d2   u2d3   u2d4   u2d5
mounted         0      0      0      0      0
vol3           u1d6   u1d7   u1d8   u1d9
mounted         0      0      0      0
vol4           u2d6   u2d7   u2d8   u2d9
mounted         0      0      0      0
```

All drives should show a status of `0`. See TABLE 5-1 for definitions of drive status codes.

Checking Data Parity

Caution – It can take up to several hours for the parity check once the `vol verify` command is executed. Execution of this command might affect system performance, depending on system activity and the verification rate selected.

- **Use the `vol verify` command to perform a parity check of the drives.**

```
:/:<7> vol verify volume name
```

You can also use the `fix` and `rate` options:

```
:/:<7> vol verify volume name [fix] [rate <1-8>]
```

Where:

- `fix` recalculates and rewrites the parity block if a mismatch is detected.
- rate specifies the speed with 1= slowest and 8 = fastest.

Note – The `vol` command is not re-entrant. Other `vol` commands cannot run on the array or partner group until the `vol verify` operation has completed.

Note – It is a good practice to run `vol verify` before recycling backup tapes to be sure the image is correct before over-writing previous images.

Checking Drive Temperature

- **Use the `fru stat` command on the array to check disk drive temperatures.**

```
:/:<43> fru stat

CTLR    STATUS   STATE       ROLE        PARTNER    TEMP
------  -------  ----------  ----------  -------    ----
u1ctr   ready    enabled     master      u2ctr      31.0
u2ctr   ready    enabled     alt master  u1ctr      30.5

...
```

Note – A warning message will appear in the array `syslog` file if a disk drive reaches 65 degrees C. The array automatically starts spinning down an individual drive if the drive's temperature reaches 75 degrees C.

Disk Drive LEDs

LEDs at the top of each disk drive indicate drive activity and status. These LEDs appear in the front cover on the unit. TABLE 5-2 lists the possible drive LED states and a description for each state.

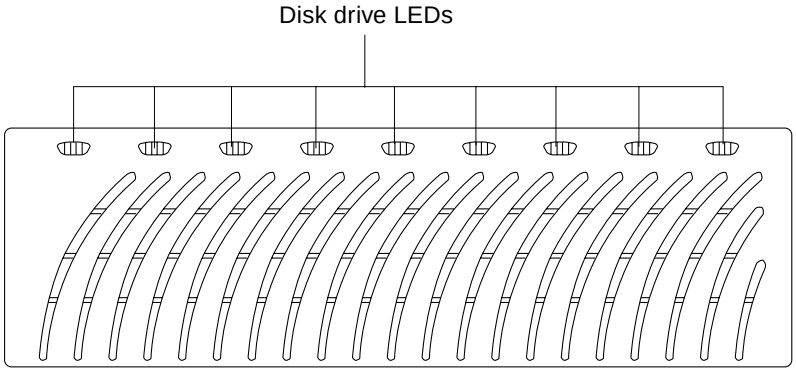

FIGURE 5-1 Disk Drive LEDs (Viewed Through Front Cover)

TABLE 5-2 Disk Drive LED Descriptions

Drive Activity (Green)	Drive Status (Amber)	Descriptions
Off	Off	Drive not installed (not recognized)
Slow blink	Off	Drive is spinning up or down
Solid	Off	Drive OK, idle

TABLE 5-2 Disk Drive LED Descriptions

Drive Activity (Green)	Drive Status (Amber)	Descriptions
Flashing	Off	Drive OK, activity
Off	Solid	Drive reconstruct/firmware download in progress
Off	Slow blink	Drive failure; OK to replace drive

Note – Even if the LED indicates a drive failure, always verify the FRU status using the CLI before replacing the drive. Refer to “Checking FRU Status” on page 35 for instructions.

Repairing Disk Drives

Caution – Replace only one disk drive in a array at a time to ensure that no data is lost. Before replacing another disk drive in the same array, complete any volume reconstructions before and ensure that the disk drive is fully functional and in operation.

By default the array automatically spins up and reenables a replaced disk drive, then automatically reconstructs the data from the parity or hot spare disk drives. The disk drive spinup takes about 30 seconds. Reconstruction of the data on the disk drive can take up to several hours depending on system activity.

Note – For the array to automatically reconstruct drive data, the array must remain powered on while a disk is replaced.

Removing and Replacing a Disk Drive

1. **Observe static electricity precautions.**

   See “Static Electricity Precautions” on page 5

2. **Remove the front panel by pressing in on the side latches and pulling the cover forward. See** FIGURE 5-2**.**

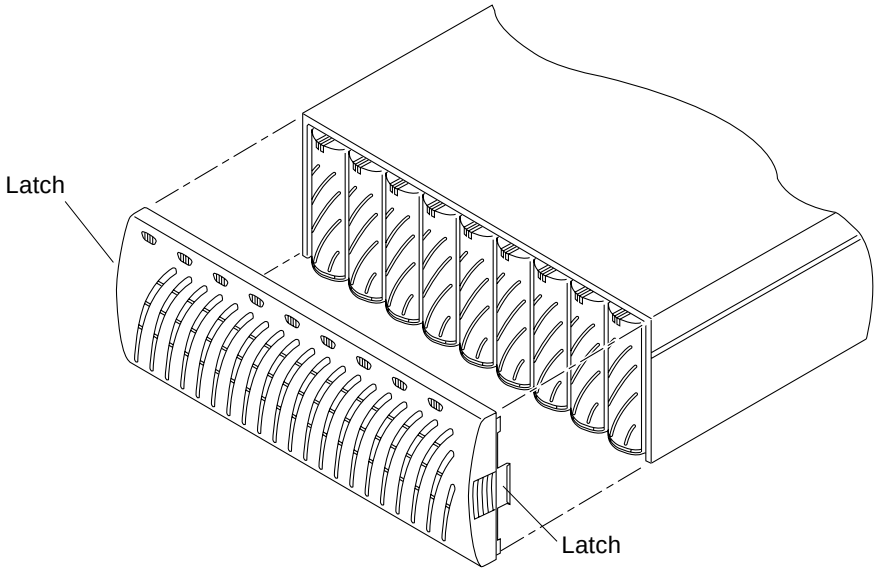

FIGURE 5-2 Removing the Front Panel

3. **Locate the disk drive that needs to be replaced.**

   Disk drives are numbered from 1 to 9 starting on the left side of the array.

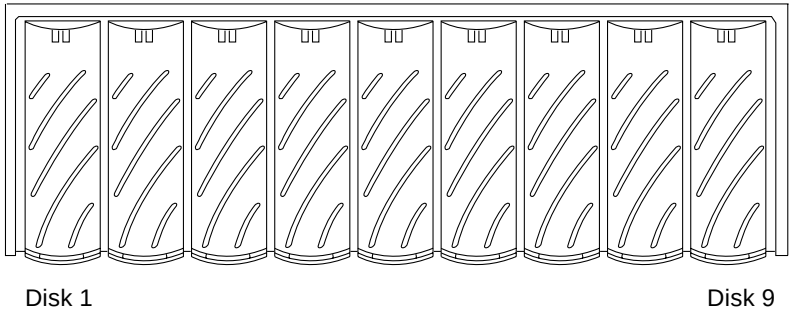

FIGURE 5-3 Disk Drive Numbering

4. **Use a coin or small screwdriver to press in and release the drive latch handle.**

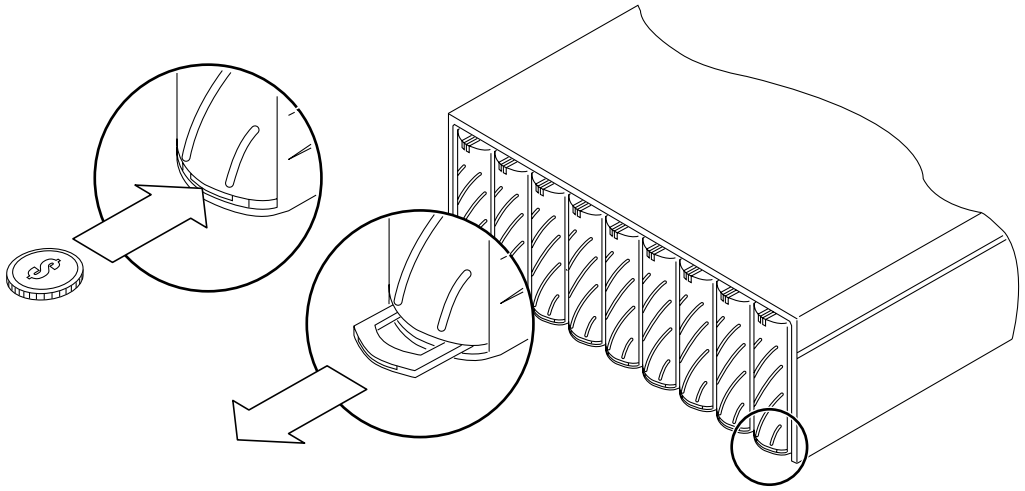

FIGURE 5-4 Releasing the Latch Handle

5. **Use the latch handle to slowly pull the disk drive out 1 inch (2.5 cm).**

   Wait 30 seconds, and then pull the drive out completely. This gives the disk drive time to spin down.

6. **Remove the disk drive from the array. See** FIGURE 5-5**.**

   Push in the latch handle on the removed disk drive to protect it from damage.

Caution – Any disk drive that is removed must be replaced within 30 minutes or the Sun StorEdge T3+ array and all attached arrays will automatically shut down and power off.

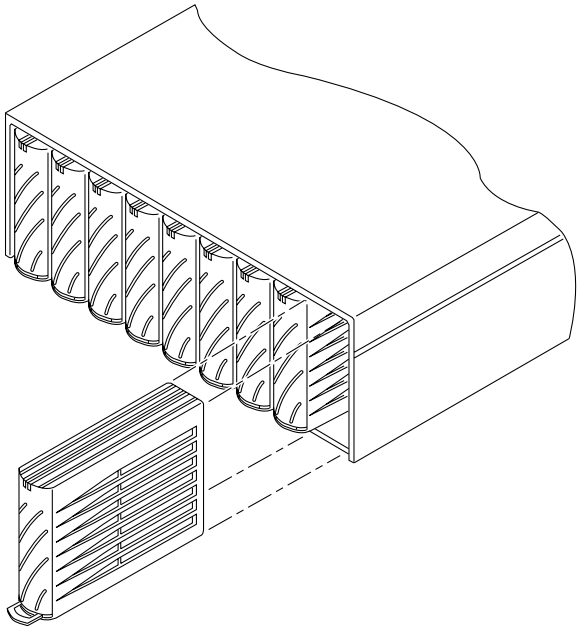

FIGURE 5-5 Removing a Disk Drive

7. **Release the latch handle on the disk drive to be installed.**

8. **Insert the new disk drive gently on the middle of the rails and push it in until it is seated with the centerplane connector.**

   Use a coin or small screwdriver to press in and lock the latch handle.

9. **Replace the front panel.**

Note – Replace the front panel for the array to meet FCC compliance requirements.

---

10. **Type `fru list` u*n*d*n* to verify the firmware revision of the new disk drive, where:**

    - un is the unit (u) number (*n*)
    - dn is the drive (d) number(*n*).

    See “Check the drive status to ensure that the reconstruction of the replaced drive FRU has completed.” on page 70,” for instructions, if necessary.

Rebuilding a Replaced Drive

A replaced drive should begin to rebuild itself automatically.

Note – If a standby drive is configured, data is not copied back from the hot spare to a newly replaced data drive until the reconstruction of data to the hot spare from parity is completed. This means that you might not see any activity lights immediately after replacing a drive.

If automatic reconstruction does not start or fails, begin the rebuild of the replaced drive FRU manually as follows:

1. **On the host, use the** `telnet` **command with the array name (or IP address) to connect to the array.**

```
mngt_host# telnet array name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

Telnet session (129.150.47.101)
```

2. **Log in to the array by typing** `root` **and the supervisor password at the prompts.**

3. **On the array, type**:

```
:/:<34> vol recon volume-name [from_standby]
```

4. **Start a second Telnet session with the array to check rebuild progress.**

5. **Check the rebuild progress.**

   Use the information in the `PERCENT` column and the `TIME` column, which shows the elapsed time, for estimating when the volume will complete reconstruction.

```
:/:<35> proc list

VOLUME          CMD_REF PERCENT    TIME COMMAND
v1                20241      23    0:09 vol recon
```

Note – If all power is removed from the array while the drive is being reconstructed, the reconstruction process restarts at the beginning when power is restored.

6. **Check the drive status to ensure that the reconstruction of the replaced drive FRU has completed.**

   The following example shows a standby drive configured for each volume.

```
:/:<43> fru stat
CTLR    STATUS   STATE       ROLE        PARTNER    TEMP
------  -------  ----------  ----------  -------    ----
u1ctr   ready    enabled     master      u2ctr      30.5
u2ctr   ready    enabled     alt master  u1ctr      30.5

DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      30    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      28    vol1
u1d6    ready    enabled     data disk   ready      ready      29    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      31    vol3
u2d1    ready    enabled     data disk   ready      ready      34    vol2
u2d2    ready    enabled     data disk   ready      ready      38    vol2
u2d3    ready    enabled     data disk   ready      ready      36    vol2
u2d4    ready    enabled     data disk   ready      ready      37    vol2
u2d5    ready    enabled     data disk   ready      ready      34    vol2
u2d6    ready    enabled     data disk   ready      ready      35    vol4
u2d7    ready    enabled     data disk   ready      ready      35    vol4
u2d8    ready    enabled     data disk   ready      ready      40    vol4
u2d9    ready    enabled     data disk   ready      ready      36    vol4

LOOP    STATUS   STATE       MODE        CABLE1     CABLE2     TEMP
------  -------  ----------  -------     ---------  ---------  ----
u2l1    ready    enabled     master      installed  -          29.5
u2l2    ready    enabled     slave       installed  -          31.0
u1l1    ready    enabled     master      -          installed  29.5
u1l2    ready    enabled     slave       -          installed  30.5

POWER   STATUS   STATE       SOURCE  OUTPUT  BATTERY  TEMP    FAN1    FAN2
------  -------  ---------   ------  ------  -------  ------  ------  ------
u1pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u1pcu2  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu2  ready    enabled     line    normal  normal   normal  normal  normal
```

Upgrading Disk Drive Firmware

The latest disk drive firmware versions are located on the SunSolve web site:

```
http://sunsolve.sun.com
```

During a disk drive firmware download, the functionality of the array is limited. To avoid system problems, verify that:

- A current backup copy of array data exists.
- The data path to the host has been quiesced. There must not be any I/O activity during the disk drive firmware download.
- The Ethernet connection is not being used for any other operation during this procedure.

Caution – If a host-mounted utility program is actively polling, problems might occur during the firmware download. Disable the polling utility during this procedure to avoid problems.

- No unnecessary command line program interaction with the array is performed during the disk drive firmware download.

Note – The disk firmware download takes approximately 20 minutes for 9 drives. Do not attempt to interrupt the download or perform other command-line functions during the process. The command prompt reappears after the download process has completed.

To upgrade the firmware:

1. **Use `ftp` to transfer the firmware to the `array root` directory in binary mode.**

   See “Establishing an FTP Session” on page 12 for additional information.

   Note – The file name of files being transferred to the local disk must be 12 characters or less in size and start with an alphabetic character (not numeric).

2. **Establish a Telnet connection to the array.**

   See “Establishing a Telnet Session” on page 9.

3. **Verify that all disk drives are in an optimal state as follows:**

   a. **Use the** `fru stat` **command to confirm that all disks are ready and enabled.**

   b. **Use the** `vol stat` **command to confirm that all disks that are configured into volumes are in an optimal state, reported as drive state** `0`**.**

   If either of these commands display drive issues, correct problems before proceeding with the firmware download.

4. **Use the** `proc list` **command to verify that there are no volume operations in progress.**

   Allow a volume operation in progress to complete before proceeding with the firmware download.

5. **Use the** `refresh -s` **command to verify that there are no battery refresh operations in progress.**

   Allow a battery refresh in progress to complete before proceeding with the firmware download.

6. **Unmount the array volume(s) from the host to ensure there is no host I/O activity.**

```
# unmount /t3 filesystem name
```

7. **Unmount internal array volume(s).**

```
:/:<1> vol unmount volume name
```

8. **Install the firmware using the** `disk download` **command.**

```
:/:<2> disk download u1d1-9 filename
```

   The *filename* is the file name of the disk drive firmware image that was transferred by FTP to the array in Step 1.

Caution – If the array is configured with different manufacturers types of disk drives, the `disk` command can download firmware for only one manufacturers drive type at a time. Verify that the download was successful using either the CLI.

9. **Use the** `fru list` **command to verify that the firmware download was successful.**

   The current drive firmware level is displayed in the `fru list` output.

10. **Use the `reset` command to reboot the Sun StorEdge T3+ array after all drives have been upgraded.**

```
:/:<3> reset
```

11. **After the array is back online, log in to the array and verify that all FRU states are optimal as follows:**

    a. **Use the `fru stat` command to confirm that all drives are ready and enabled.**

    b. **Use the `fru list` command to display the current drive model number and firmware version.**

    c. **Use the `vol stat` command to display drive states.**

    All drives must report a drive state of `0` for optimal condition.

12. **Remount the volume(s) on the array.**

```
:/:<4> vol mount volume name(s)
```

CHAPTER 6

Interconnect Card Assemblies

This chapter describes how to monitor and replace the interconnect card, and upgrade firmware. The chapter contains the following sections:

Interconnect Card LEDs

Each of the interconnect cards has a status LED for each interconnect cable. TABLE 6-1 lists the possible interconnect card status LED states with descriptions of each state.

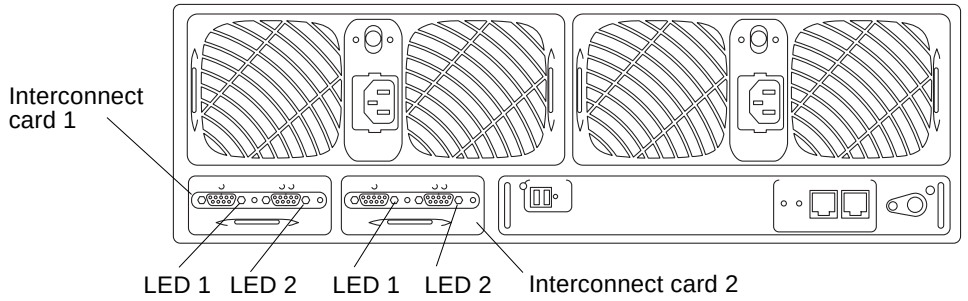

FIGURE 6-1 Interconnect Card LEDs

TABLE 6-1 Interconnect Card LED Descriptions

Interconnect Card Status LED (Green or Amber)	Description
Off	Interconnect card not installed (not recognized)
Green—solid	Interconnect card OK Cable OK (if present)
Green—slow blink	Interconnect card OK, possible communication problem with other cards. Cable may be bad, OK to replace cable
Amber—solid	Interconnect card firmware download in progress
Amber—slow blink	Interconnect card failure, OK to replace interconnect card

Note – Even if the LED indicates an interconnect card failure, always verify the FRU status using the CLI before replacing the interconnect card. Refer to “Checking FRU Status” on page 35 for instructions.

Removing and Replacing an Interconnect Card

Caution – Use the interconnect cables only for cabling Sun StorEdge T3+ arrays together using the interconnect card connectors. Do *not* use these cables for any other FC-AL connection.

Caution – The interconnect card is extremely sensitive to static electricity. Use proper antistatic wrist straps and antistatic procedures when handling any FRU.

Caution – Replace one interconnect card at a time. Pulling both interconnect cards at one time could cause a system shutdown. Follow the procedure as described to ensure that there is no interruption in system operation or loss of data.

To prevent interruption of the data host system operation during interconnect card replacement, ensure that:

- In a single controller unit configuration, remove only the failed interconnect card. Leave the second interconnect card intact in the array.
- In a partner group, remove the interconnect cable only from the failed interconnect card. Leave the interconnect cable attached to the working interconnect card.

To replace an interconnect card:

1.Ensure that the interconnect card to be replaced is showing failure status.

Refer to FIGURE 6-1.

2.Observe static electricity precautions.

See “Static Electricity Precautions” on page 5.

3.Remove the interconnect cable from the failed interconnect card only.

Note – If a single controller-unit configuration, ignore this step and proceed to Step 4.

Mark the connector with either 1 or 2.

4.Unlock the failed interconnect card by pushing in on the latch handle.

Use a coin or small screwdriver to press in and release the latch handle.

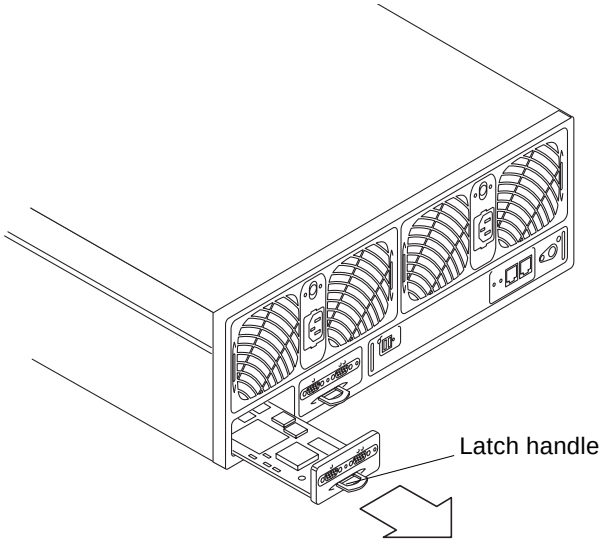

FIGURE 6-2 Removing the Interconnect Card

5.Pull the interconnect card out using the latch handle.

Caution – The interconnect card that is removed must be replaced within 30 minutes or the Sun StorEdge T3+ array and all attached arrays will automatically shut down and power off.

6.Insert the new interconnect card, making sure that the card sits on the frame.

7.Lock the new interconnect card in place by pushing in the latch handle.

Use a coin or small screwdriver to press in and secure the latch handle.

8.Reconnect the interconnect cable to the interconnect card.

9.Verify that the LEDs on the interconnect card show that the interconnect card has initialized properly.

10. Verify the status of the interconnect card using the CLI.

Refer to “Checking FRU Status” on page 35 for instructions.

11. **Type** `lpc version` **to view and verify the firmware level of the new interconnect card.**

    See “Upgrading Interconnect Card Firmware” on page 79” for instructions, if necessary.

Upgrading Interconnect Card Firmware

The interconnect card firmware is stored in the FLASH memory device on the interconnect card. The array can be operational during the interconnect card firmware upgrade.

The firmware upgrade procedures that follow must be done through the Ethernet connection. The latest firmware versions are located on the SunSolve web site:

```
http://sunsolve.sun.com
```

- The firmware must be resident on the host for this operation.
- The Sun StorEdge T3+ arrays must have a supervisor password prior to attempting this procedure.

To upgrade the firmware, see the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

CHAPTER 7

Power and Cooling Unit Assemblies

This chapter describes how to replace the power and cooling unit and monitor the UPS battery. The chapter contains the following sections:

Power and Cooling Unit

The power and cooling unit has two active power sources: standby and primary power. Standby power, which is used to power the micro controller on the interconnect card, is activated when AC power is present. Primary power, which is used to power all remaining circuits and disk drives, is activated when AC or battery power is present and the power switch is on.

Each power and cooling unit has a power switch in the rear upper center of the unit. Turning off the power on a power and cooling unit affects only that power and cooling unit. Therefore, to power off all primary power to the unit, both power switches on both power and cooling units must be turned off. After the switches are turned off, system primary power will not actually turn off until the controller has performed an orderly shutdown, including writing any data cache to disk. This process can take up to two minutes.

Separate power cords are used for the connector on each power and cooling unit to provide redundant cabling. The power cords need to be connected to separate AC power sources for full redundancy.

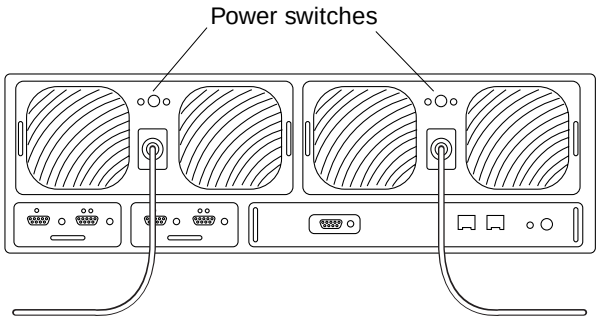

FIGURE 7-1 Power Cords Connected to the Power and Cooling Units

Caution – Do not handle the power and cooling unit when the power cord is connected. Line voltages are present within the power and cooling unit when the power cord is connected even if the power switch is off.

At the rear of the power and cooling unit is a recessed PC card connector. Do not touch this connector or allow any metal object to touch it. The power and cooling unit contains the UPS battery backup.

Note – The batteries in the power and cooling units recharge after powering on the array. If the batteries are less than fully charged, `fru stat` output displays batteries in a “fault” condition, and write-behind cache is disabled until the batteries are charged. The system can take several hours to determine the health of the batteries after the system is turned back on. Batteries reflect a non optimal state after power loss events and also after turning off power switches.

Power and Cooling Unit LEDs

Each of the power and cooling units has an AC LED and a power supply (PS) LED. TABLE 7-1 lists the possible conditions of these LEDs with a description of each state.

Power and Cooling Unit LEDs

Each power and cooling unit has an AC LED and a power-supply (PS) LED. TABLE 7-1 lists the possible conditions of these LEDs and describes each state.

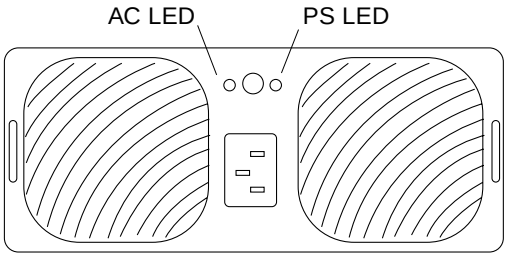

FIGURE 7-2 Power and Cooling Unit LEDs

TABLE 7-1 Power and Cooling Unit LED Descriptions

AC LED (Green or Amber)	PS LED (Green or Amber)	Description
Off	Off	• Power is off • No AC input
Amber	Off	• Power is off • Power switch turned off • AC power is available
Green	Off	Occurs when array is shut down: • PCU disabled • AC power is available
Green	Green	Normal operating state: • PCU receiving AC power • Power switch is turned on • AC power is available

TABLE 7-1 Power and Cooling Unit LED Descriptions *(Continued)*

AC LED (Green or Amber)	PS LED (Green or Amber)	Description
Amber	Amber	• Switch is off. Array powers off after PCU is disabled
Green	Amber	Indicates one or more of following: • Over-temperature condition; PCU disabled • DC power not available; PCU disabled • Both fans fault; PCU disabled • Battery on refresh cycle
Green	Blinking green	• Battery not ready; charging
Green	Blinking amber	Indicates one or more of following: • PCU disabled • One fan fault • Battery hold-time low; PCU remains enabled • Battery out of warranty; PCU remains enabled • Battery life-span failure; PCU remains enabled Note—Verify a power and cooling unit failure using the CLI or Component Manager.

Note – Even if the LED indicates a power cooling unit failure, always verify the FRU status using the CLI before replacing the power cooling unit. Refer to “Checking FRU Status” on page 35 for instructions.

Removing and Replacing a Power and Cooling Unit

Caution – To ensure correct airflow for system cooling, both power and cooling units must be in the installed position for normal operation. A failed power and cooling unit should be removed only when a replacement power and cooling unit is available to be inserted.

Caution – Replace only one power and cooling unit at a time to prevent system interruption.

To replace a power and cooling unit:

1. **Observe static electricity precautions.**

   See “Static Electricity Precautions” on page 5

2. **Power off the power and cooling unit by pressing the power switch (**FIGURE 7-1**).**

   Make sure that the AC LED is amber and the PS LED is off (FIGURE 7-2).

3. **Disconnect the power cord from the AC outlet.**

4. **Disconnect the power cord from the power and cooling unit connector by squeezing both sides of the connector and pulling straight out (**FIGURE 7-1**).**

5. **Unlock the power and cooling unit by using a coin or small screwdriver to push in and release the two latch handles** FIGURE 7-3**).**

6. **Pull the power and cooling unit out of the array.**

   Put one index finger through each of the latch handles. With your thumbs on the top of the chassis for support, pry the power and cooling unit out of its connectors with an upward rotation. Once it is out approximately 1 inch (2.5 cm), the unit will be free to slide out of the frame on its rails.

Caution – Any power and cooling unit that is removed must be replaced within 30 minutes or the Sun StorEdge T3+ array and all attached arrays automatically shut down and power off.

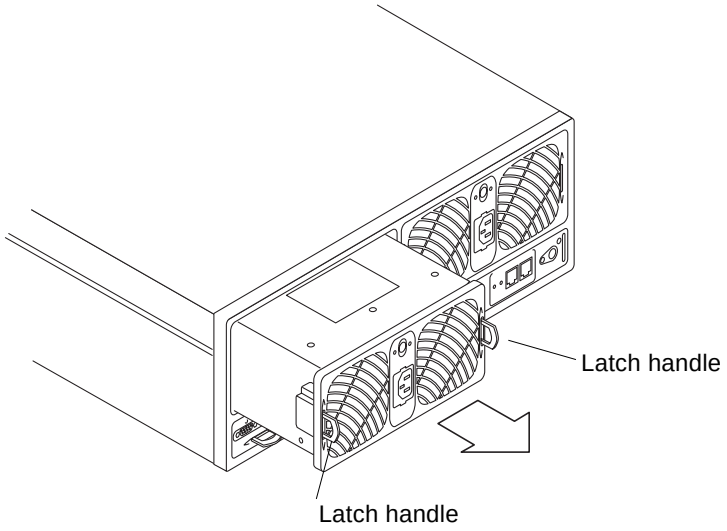

FIGURE 7-3 Removing the Power and Cooling Unit

7. **Insert the new power and cooling unit.**
8. **Lock the new power and cooling unit by pushing in both latch handles.**
9. **Insert the power cord into the power and cooling unit connector.**
10. **Connect the power cord into the AC outlet.**

    Verify that the AC LED on the power and cooling unit is amber, indicating that AC power is present.
11. **Push the power and cooling unit power switch on.**

    Verify that both LEDs on the power and cooling unit are green, indicating that the unit is receiving power.
12. **Verify the status of the power and cooling unit using the CLI.**

    Refer to "Checking FRU Status" on page 35 for instructions.

Note – After installing the new power and cooling unit, the batteries will take some time to recharge.

UPS Battery

The uninterruptible power supply (UPS) battery is located within the power and cooling unit. The battery provides backup in case of a complete AC power failure and sustains power to the array long enough to flush cache data to the drives.

When a Sun StorEdge T3+ array is first powered up, write-behind caching is disabled (cache runs in write-through mode) for a short time. The write-behind caching is disabled during cold boots (even if AC power has not been removed from the array) as the firmware attempts to determine the condition of the internal PCU batteries. Once the system determines that the batteries are in an optimal state, system cache mode returns to write-behind. After a power down, a array re-enables write-behind cache mode in approximately two hours.

During a power failure, if the battery is flushing cache data to the drives, battery power becomes depleted. Once AC power is available, the battery recharges. While the battery is recharging, write-behind cache mode is disabled and write-through cache mode is enabled until the battery is fully recharged. The battery recharge could take up to 12 hours, depending on the length of the power outage and the amount of cache data that was flushed to the drives.

Note – The batteries in the power and cooling units recharge after powering on the array. If the batteries are less than fully charged, `fru stat` output will display batteries in a fault condition, and write-behind cache is disabled until the batteries are charged.

Checking the Battery

1. **On the host, use the `telnet` command with the array name (or IP address) to connect to the array.**

```
mngt_host# telnet array-name
Trying 123.123.123.101...
Connected to 123.123.123.101.
Escape character is '^]'.

pSOSystem (123.123.123.101)
```

2. **Log in to the array by typing `root` and the supervisor password at the prompts.**

3. **Use the** `id read` **command to display battery life related information.**

   (Unit number *n* = 1 or 2; power cooling unit number *n* = 1 or 2.)

```
:/: id read unpcun
Revision              : 0000
Manufacture Week      : 00281999
Battery Install Week  : 00412001
Battery Life Used     : 275 days, 2 hours
Battery Life Span     : 730 days, 12 hours
Serial Number         : 001787
Battery Warranty Date: 20011119142702
Battery Internal Flag: 0x00000000
Vendor ID             : TECTROL-CAN
Model ID              : 300-1454-01(50)
```

4. **Use the** `refresh -s` **command to check the status of a battery refresh cycle.**

   The following examples show a battery refresh in progress and a normal battery status (no refresh cycle):

```
:/: refresh -s

PCU1                      PCU2
-------------------------------------------------------------------
U1                        Completed                 Recharging

Current Time              Fri May 26 18:32:07 GMT 2002
Start Time                Thu May 25 20:31:19 GMT 2002
Last Refresh              Thu May 11 20:22:53 GMT 2002
Next Refresh              Thu Jun 08 20:31:19 GMT 2002
Total time elapsed:       22 hours, 0 minutes, 48 seconds.
```

```
:/: refresh -s
No battery refreshing Task is currently running.

                          PCU1                            PCU2
-------------------------------------------------------------------
U1                        Normal                          Normal
U2                        Normal                          Normal

Current Time              Wed Aug 21 16:45:36 GMT 2002
```

Battery Maintenance

The battery refresh cycle occurs automatically once every 28 days. The battery refresh cycle is sequential, ensuring that only one battery in a unit is refreshed at a time. The refresh cycle consists of a 6 minute discharge period, followed by a recharge period of 6 to 12 hours.

The refresh cycle verifies the health of the battery. During the refresh, if a problem is detected with the battery, future refresh operations are suspended until the problem is fixed. When refresh is suspended, battery write-behind caching is turned off automatically as a safety precaution.

The `syslog` file indicates battery refresh operation in progress. Use the `refresh -s` command to view an active refresh operation. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for more information on this command. Refresh cycle time is controlled by the array's `/etc/schd.conf file`. For example, specify that a battery refresh cycle begin on January 15, 2001 at 11 p.m., the entry in the `/etc/schd.conf` file:

```
:/: cat /etc/schd.conf
BAT_BEG 1-15-2001,23-00-00
BAT_CYC 28
```

You can tune the `/etc/schd.conf` file to specify the interval between battery refresh cycles and initiate a refresh on a particular day. To specify beginning a battery refresh cycles at a particular time, edit the BAT_BEG *MM-DD-YYYY,hh-mm-ss* value in the `/etc/schd.conf` file. Where:

- *MM* is the month number (January = 1)
- *DD* is the day number
- *YYYY* is the year
- *hh* is the hour using a 24 hour clock (6pm = 18)
- *mm* is the minute
- *ss* is the second (this element is optional)

Caution – The battery service life dependents on a battery refresh cycle of 28 days. Altering this time span can decrease battery life and should only be done as directed by Sun representatives.

Note that the next refresh start time is always calculated from the start time of the previous refresh cycle. If a user manually starts a refresh cycle, then the next refresh depends on the starting time of the manually activated refresh cycle.

Note – If a controller failover occurs, the scheduler daemon starts and behave as it does during a normal system boot. The scheduler reads the `schd.log` file, and based on `schd.conf` file, begins the next refresh process. If during the discharge period (6 minutes) or recharge period (6 to 12 hours) a controller failover occurs, the current refresh process is killed and the next refresh cycle starts at the scheduled refresh time based on the `schd.conf` file. Consequently, the refresh cycles begins as scheduled previously.

The battery service life is 2 years. When the battery approaches its end of life, warning messages are sent to the `syslog` file. The first message is sent 45 days before the end of life, followed by a warning message every five days thereafter. The power and cooling unit *must* be replaced within forty-five days of receiving the first warning message. The warning message indicates which power and cooling unit needs to be replaced. After the battery service life expires, the cache is forced to write-through mode.

Removing and Replacing the UPS Battery

This section covers a cold-swap procedure.

Note – Eventhough the T3+ documentation targeted for the customers suggests replacing the PCU to replace the battery, trained Sun Field personnel, who have access to the Sun StorEdge Field Service Manual, may follow the procedure shown below to replace only the battery inside the PCU.

Remove the UPS Battery

1. **Remove the PCU from the array.**

   See “Removing and Replacing a Power and Cooling Unit” on page 85.

Caution – Any power and cooling unit that is removed must be replaced within 30 minutes or the Sun StorEdge T3+ array and all attached arrays automatically shut down and power off.

2. **Turn the PCU over such that the bottom of the unit is facing up as shown in** FIGURE 7-4**.**

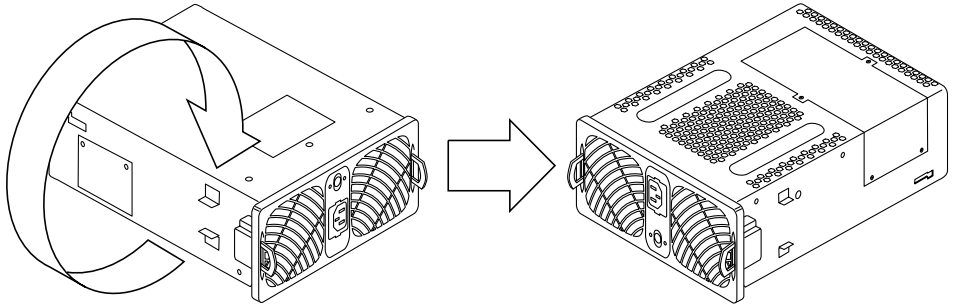

FIGURE 7-4 Turning the PCU upside down

3. **Remove the four Phillips screws from the panel on the bottom and side of the PCU as shown in** FIGURE 7-5**.**

   Use care in removing the screws so they do not fall into the vent holes of the PCU.

Note – The battery is attached to the bottom panel of the PCU. When removing the bottom panel, do not attempt to remove it completely as the battery is still connected to the unit.

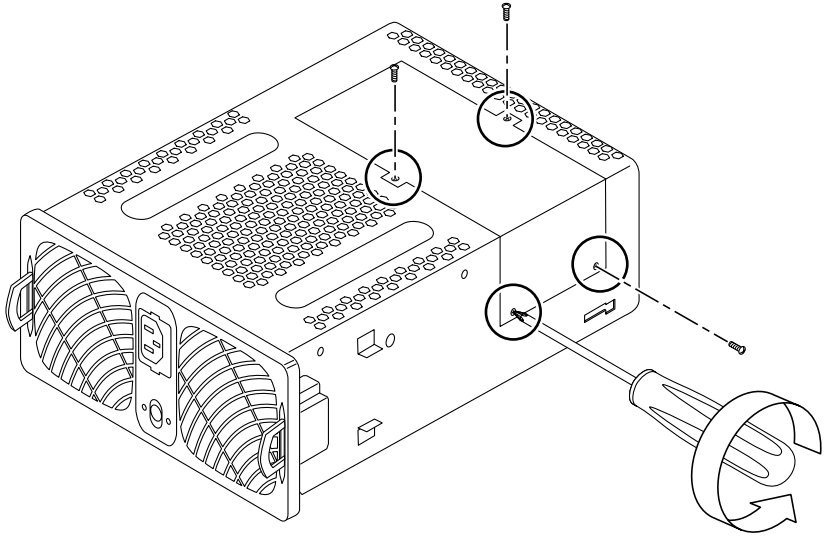

FIGURE 7-5 Removing the Screws from the PCU Bottom Panel

4. **Slide the bottom panel off the unit slightly, enough to expose the battery connector as shown in** FIGURE 7-6 **and** FIGURE 7-7**.**

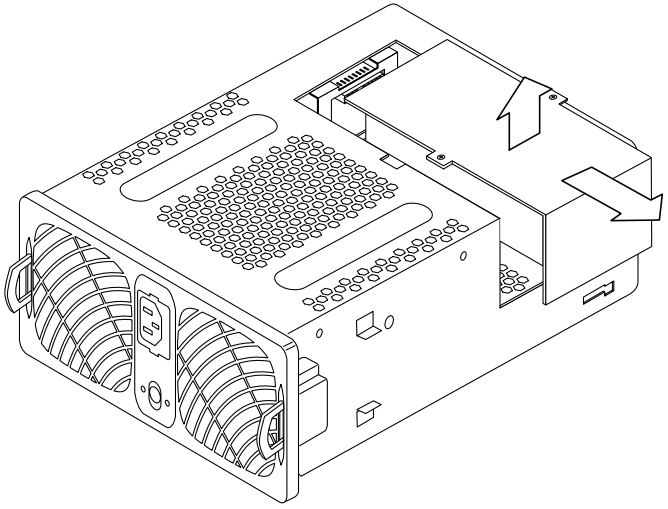

FIGURE 7-6 Lifting the PCU Bottom Panel and Battery Slightly Away from the Unit

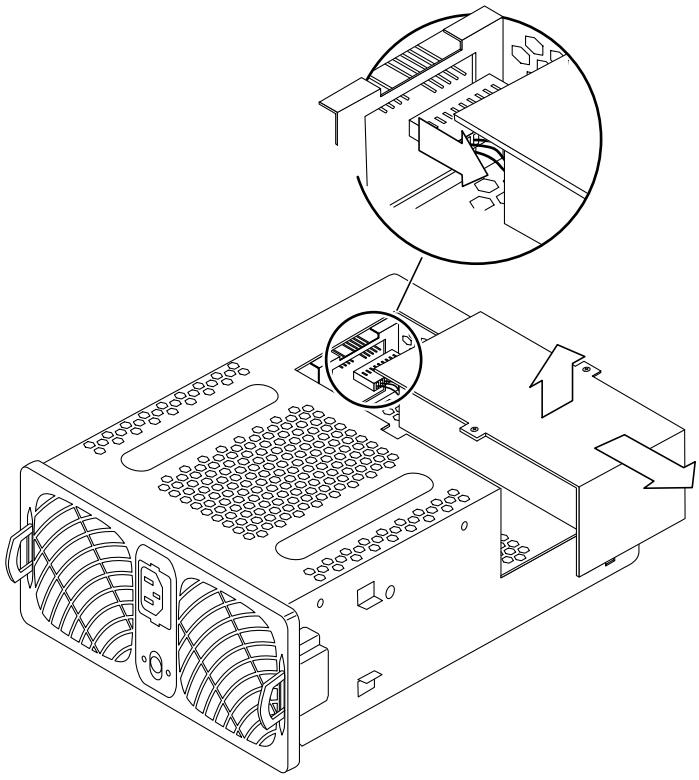

FIGURE 7-7 The Battery Connector Details Inside the PCU

5. **Remove the battery connector by pulling on if firmly straight out from the connector inside the PCU.**

6. **Lift the bottom panel with the battery away from the unit and set it aside as shown in** FIGURE 7-8**.**

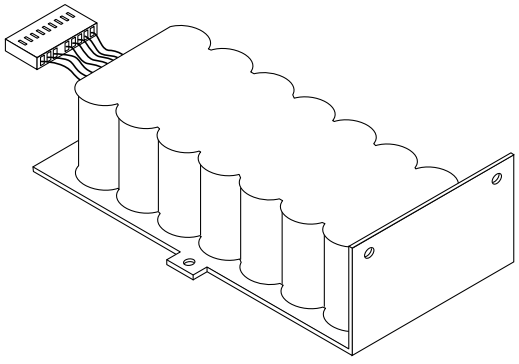

FIGURE 7-8 UPS Battery Setting Right Side Up

Replace the UPS Battery

1. **Connect the replacement battery to the battery connector of the PCU.**

   See FIGURE 7-7.

   Firmly push the connector all the way into the PCU battery connector. There is no indication, such as a mechanical click, that indicates that it is fully inserted.

2. **Seat the battery pack in the PCU such that the bottom panel is flush with the edges of the PCU.**

   See FIGURE 7-5.

3. **Replace the four Phillips screws and secure the bottom panel to the PCU.**

4. **Replace the PCU in the array and prepare the PCU for return to service as described in “Removing and Replacing a Power and Cooling Unit” on page 85.**

5. **Reset the date by typing** `.bat -n u(x) pcu(y)` **from the T3+ CLI prompt, where** `u(x)` **is the unit number and** `pcu(y)` **is the location number associated with the PCU that was just installed.**

   This command will zero out the “Battery Warranty Date” field and set the “Battery Install Week” field according to the T3 date setting. Additionally, this command will zero out the “Battery Internal Flag” field if it was set to 1 to indicate low battery.

6. **Type** `.id write busage u(x)pcu(y)0` **from the T3+ CLI prompt**

   This command calculates the “Battery Warranty Date” and “Battery Life Used.” To verify this, you may type `id read u(x)pcu(y)`.

CHAPTER 8

Diagnosing and Correcting FC-AL Loop Problems

This chapter describes how to diagnose and correct back-end FC-AL drive loop problems with the array. It contains the following sections:

There are several failure conditions within the back-end loop that do not appear as a `failed FRU` status. These kind of failures can only be diagnosed by collecting data from various sources within the system such as, `iostat` performance data, CLI status commands, Storage Automated Diagnostic Environment, Storage Automated Diagnostic Environment message monitoring, the Sun StorEdge T3+ array `syslog`, and the FC-AL connected host messages file. Data from these sources is used to determine the most likely failed FRU within the Sun StorEdge T3+ array system.

Overview

The procedures in this chapter assume that the person servicing the equipment has been trained on the product and that the required service manuals are available.

A serial maintenance cable kit must be available (part number 370-4119), along with a terminal or host port connection.

Note – In order to collect the information required to diagnose back-end FC-AL loop problems, several of the engineering-only "dot" commands must be used. Only the status options of these dot commands are used.

Diagnosing and correcting back-end FC-AL loop problems can take up to five steps:

1. **Determine that there has been a failure in the back-end drive loop.**

   Diagnosing the problem requires that you analyze the collected data and make a determination of which is the most likely failed FRU from the data available. This procedure is described in “Diagnosing an FC-AL Loop” on page 105.

   Once you identify a suspected FRU, use one or more of the following steps to isolate and then replace the failed FRU.

2. **Isolate, replace, and verify the interconnect cards and/or the loop cable.**

   Interconnect cards (sometimes referred to as *unit interconnect cards* or *UIC*) can be replaced without affecting the online operation of the product, though there may be some performance impact. See “Interconnect Card Replacement Procedure” on page 115.

3. **Isolate, replace, and verify the RAID controllers.**

   Replacing RAID controllers causes a LUN/controller path failover. This failover might require some kind of manual procedure by the customer to continue running and it might affect the overall system performance. See “RAID Controller Replacement Procedure” on page 116.

4. **Isolate, replace, and verify the FC-AL disk drives.**

   Perform the this step only if steps 2 and 3 fail to resolve the problem. To run the loop diagnostics to identify a failed drive FRU, the Sun StorEdge T3+ array must be removed from operation. Removing the array is highly disruptive to the customer. See “Off-Line Drive Diagnostics and Replacement” on page 117.

5. **Replace and verify the chassis and mid-plane.**

   If by the end of step 4 there is still a problem the chassis and mid-plane will need to be replaced. Perform this step only if steps 2, 3, and 4 fail to resolve the problem. This is highly disruptive to the customer. See “Chassis Replacement Procedure” on page 123 and “Replacing the Chassis/Backplane Assembly” on page 126.”

Normal Status

The normal configuration information can be determined by using the following CLI commands and interpreting the results.

- `fru stat` (see “The `fru stat` Command” on page 98)
- `vol mode` (see “The `vol mode` Command” on page 99)
- `port listmap` (see “The `port listmap` Command” on page 100)
- `.loop stat` (see “The `.loop stat` Command” on page 101)

- `.disk pathstat` (see "The `.disk pathstat` Command" on page 101)
- `.disk linkstat` (see "The `.disk linkstat` Command" on page 103)

The examples that follow show a Sun StorEdge T3+ array in a redundant partner group configuration, with no failed FRUs.

The `fru stat` Command

The `fru stat` command returns the current condition of both disk ports (port 1 and port 2), as well as the status of the interconnect cards. If there are loop problems, this might indicate certain disk ports have a status other than `ready`, or the loop cards with a status other than `ready` or `enabled`.

```
:/:<43>fru stat

CTLR    STATUS   STATE       ROLE        PARTNER    TEMP
------  -------  ----------  ----------  -------    ----
u1ctr   ready    enabled     master      u2ctr      30.5
u2ctr   ready    enabled     alt master  u1ctr      30.0
DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      29    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      28    vol1
u1d6    ready    enabled     data disk   ready      ready      29    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      31    vol3
u2d1    ready    enabled     data disk   ready      ready      33    vol2
u2d2    ready    enabled     data disk   ready      ready      37    vol2
u2d3    ready    enabled     data disk   ready      ready      35    vol2
u2d4    ready    enabled     data disk   ready      ready      37    vol2
u2d5    ready    enabled     data disk   ready      ready      34    vol2
u2d6    ready    enabled     data disk   ready      ready      35    vol4
u2d7    ready    enabled     data disk   ready      ready      35    vol4
u2d8    ready    enabled     data disk   ready      ready      40    vol4
u2d9    ready    enabled     data disk   ready      ready      36    vol4

LOOP    STATUS   STATE       MODE        CABLE1     CABLE2     TEMP
------  -------  ----------  -------     ---------  ---------  ----
u2l1    ready    enabled     master      installed  -          29.0
u2l2    ready    enabled     slave       installed  -          30.5
u1l1    ready    enabled     master      -          installed  29.5
u1l2    ready    enabled     slave       -          installed  30.0

POWER   STATUS   STATE       SOURCE  OUTPUT  BATTERY  TEMP    FAN1    FAN2
------  -------  ---------   ------  ------  -------  ------  ------  ------
u1pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u1pcu2  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu2  ready    enabled     line    normal  normal   normal  normal  normal
```

The `vol mode` Command

The `vol mode` command returns the current cache mode. A cache status other than `writebehind`, might indicate loop problems.

CODE EXAMPLE 8-1 vol mode Command—Normal Ouputs

```
:/:<2> vol mode

volume        mounted writebehind  mirror
vol1          yes     writebehind on
vol2          yes     writebehind on
vol3          yes     writebehind on
vol4          yes     writebehind on
```

The `port listmap` Command

The `port listmap` command returns the current controller to volume path. One controller controlling all the configured volumes, might indicate loop problems.

CODE EXAMPLE 8-2 `port listmap` Command—Normal Output

```
:/:<3> port listmap

port   targetid   addr_type   lun   volume        owner   access
u1p1      1         hard        0    vol1          u1      primary
u1p1      1         hard        1    vol1          u1      primary
u1p1      1         hard        2    vol1          u1      primary
u1p1      1         hard        3    vol1          u1      primary
u1p1      1         hard        4    vol1          u1      primary
u1p1      1         hard        5    vol1          u1      primary
u1p1      1         hard        6    vol1          u1      primary
u1p1      1         hard        7    vol1          u1      primary
u1p1      1         hard        8    vol1          u1      primary
u1p1      1         hard        9    vol1          u1      primary
u1p1      1         hard        10   vol2          u2      failover
u1p1      1         hard        11   vol2          u2      failover
u1p1      1         hard        12   vol3          u1      primary
u1p1      1         hard        13   vol3          u1      primary
u1p1      1         hard        14   vol4          u2      failover
u1p1      1         hard        15   vol4          u2      failover
u2p1      2         hard        0    vol1          u1      failover
u2p1      2         hard        1    vol1          u1      failover
u2p1      2         hard        2    vol1          u1      failover
u2p1      2         hard        3    vol1          u1      failover
u2p1      2         hard        4    vol1          u1      failover
u2p1      2         hard        5    vol1          u1      failover
u2p1      2         hard        6    vol1          u1      failover
u2p1      2         hard        7    vol1          u1      failover
u2p1      2         hard        8    vol1          u1      failover
u2p1      2         hard        9    vol1          u1      failover
u2p1      2         hard        10   vol2          u2      primary
u2p1      2         hard        11   vol2          u2      primary
u2p1      2         hard        12   vol3          u1      failover
u2p1      2         hard        13   vol3          u1      failover
u2p1      2         hard        14   vol4          u2      primary
u2p1      2         hard        15   vol4          u2      primary
```

The `.loop stat` Command

The `.loop stat` command returns the current loop configuration with regard to the electrical connections between the loop cards. A loop configuration other than the example below, might indicate loop problems.

Note – The “+” symbol represents the presence of the ISP2200 chip.

CODE EXAMPLE 8-3 `.loop stat` Command—Normal Ouput

```
:/:<4> .loop stat

Loop 1: <1+> <2+>
Loop 2: <1+2+>
```

Where:

- `<1+>` means u1d1-9 and u1ctr ISP2200 are on the loop.
- `<2+>` means u2d1-9 are u2ctr ISP2200 are on the loop.
- `<1+><2+>` means the loop is split into 2 segments.
- `<1+2+>` means u1d1-9 and u2d1-9 and u1ctr and u2ctr ISP2200s are all on the loop.
- `<1+2>` means u1d1-9 and u2d1-9 and u1ctr ISP2200 are on the loop. A disabled u2ctr would result in this configuration.
- <12+> means u1d1-9 and u2d1-9 and u2ctr ISP2x00 are on the loop. A disabled u1ctr could result in this configuration.

The `.disk pathstat` Command

The `.disk pathstat` command returns the current disk path logical configuration. A path status other than what is displayed below, might indicate loop problems.

Note – The Telnet session always runs the command through the master controller.

CODE EXAMPLE 8-4 `.disk pathstat` Command-Normal Output

```
qatest:/:<8>.loop stat
Loop 1: <1+> <2+>
Loop 2: <1+2+>
qatest:/:<9>.disk pathstat u1d1-9

DISK PPATH  APATH  CPATH  PATH_POLICY  FAIL_POLICY
-------------------------------------------------
u1d1 [0 U]  [1 U]  APATH  APATH        PATH
u1d2 [0 U]  [1 U]  APATH  APATH        PATH
u1d3 [0 U]  [1 U]  APATH  APATH        PATH
u1d4 [0 U]  [1 U]  PPATH  PPATH        PATH
u1d5 [0 U]  [1 U]  PPATH  PPATH        PATH
u1d6 [0 U]  [1 U]  PPATH  PPATH        PATH
u1d7 [0 U]  [1 U]  PPATH  PPATH        PATH
u1d8 [0 U]  [1 U]  PPATH  PPATH        PATH
u1d9 [0 U]  [1 U]  PPATH  PPATH        PATH

pass
qatest:/:<10>.disk pathstat u2d1-9

DISK PPATH  APATH  CPATH  PATH_POLICY  FAIL_POLICY
-------------------------------------------------
u2d1 [0 U]  [1 U]  APATH  APATH        PATH
u2d2 [0 U]  [1 U]  APATH  APATH        PATH
u2d3 [0 U]  [1 U]  APATH  APATH        PATH
u2d4 [0 U]  [1 U]  APATH  PPATH        PATH
u2d5 [0 U]  [1 U]  APATH  PPATH        PATH
u2d6 [0 U]  [1 U]  APATH  PPATH        PATH
u2d7 [0 U]  [1 U]  APATH  PPATH        PATH
u2d8 [0 U]  [1 U]  APATH  PPATH        PATH
u2d9 [0 U]  [1 U]  APATH  PPATH        PATH

pass
```

Where:

- `[0 U]` means Loop 1 (path_id = 0) is Up.
- `[1 U]` means Loop 2 (path_id = 1) is Up.
- `[0 D]` means Loop 1 (path_id = 0) is Down.
- `[1 D]` means Loop 2 (path_id = 1) is Down.
- `PPATH` means primary path.

- APATH means alternate path.
- CPATH means current path.
- PATH_POLICY means the preferred path (notice the 3/6 split).
- FAIL_POLICY is not supported (always PATH for path failover vs. NONE for no failover).

The .disk linkstat Command

The .disk linkstat command returns whether a device port link status register can be accessed by a controller in its current configuration. If the link status register cannot be accessed this may indicate a path problem to those disk(s) ports.

Note – The Telnet session will always run the command through the master controller. Although it is possible to connect directly to the alternate controller, it is not supported.

CODE EXAMPLE 8-5 .disk linkstat Command—Normal Ouput

```
:/:<9> .disk linkstat u1d1-9 path 0
DISK LINKFAIL LOSSSYNC LOSSSIG PROTOERR INVTXWORD INVCRC
-------------------------------------------------------
u1d1  2        16       0       0        51        0
u1d2  2        67       0       0        48        0
u1d3  2        15       0       0        41        0
u1d4  2        56       0       0        58        1
u1d5  2        40       0       0        50        0
u1d6  2        90       0       0        39        0
u1d7  2        28       0       0        51        1
u1d8  2        20       0       0        64        1
u1d9  2        20       0       0        87        0
```

The status for the command example shown below is correct for a split loop configuration.

CODE EXAMPLE 8-6 `.disk linkstat` Command—Split Loop Ouput From U1 Controller

```
:/:<24> .disk linkstat u1d1-9 path 0

DISK LINKFAIL LOSSSYNC LOSSSIG PROTOERR INVTXWORD INVCRC
-------------------------------------------------------
u1d1  0        0        0       0        30        0
u1d2  0        0        0       0        30        0
u1d3  0        0        0       0        12        0
u1d4  1        1        0       0        249       0
u1d5  0        0        0       0        30        0
u1d6  0        2        0       0        4         0
u1d7  0        0        0       0        30        0
u1d8  0        0        0       0        30        0
u1d9  0        0        0       0        30        0

pass

:/:<25> .disk linkstat u1d1-9 path 1

DISK LINKFAIL LOSSSYNC LOSSSIG PROTOERR INVTXWORD INVCRC
-------------------------------------------------------
u1d1  0        0        0       0        30        0
u1d2  0        0        0       0        30        0
u1d3  0        0        0       0        1         0
u1d4  1        2        0       0        198       0
u1d5  0        0        0       0        30        0
u1d6  0        0        0       0        1         0
u1d7  0        0        0       0        19        0
u1d8  0        0        0       0        30        0
u1d9  0        0        0       0        30        0

pass
```

CODE EXAMPLE 8-7 `.disk linkstat` Command—Split Loop Ouput From U2 Controller

```
:/:<26> .disk linkstat u2d1-9 path 0

DISK LINKFAIL LOSSSYNC LOSSSIG PROTOERR INVTXWORD INVCRC
-------------------------------------------------------
u2d1 Disk Link Status Failed
u2d2 Disk Link Status Failed
u2d3 Disk Link Status Failed
u2d4 Disk Link Status Failed
u2d5 Disk Link Status Failed
u2d6 Disk Link Status Failed
u2d7 Disk Link Status Failed
u2d8 Disk Link Status Failed
u2d9 Disk Link Status Failed

fail

:/:<27> .disk linkstat u2d1-9 path 1

DISK LINKFAIL LOSSSYNC LOSSSIG PROTOERR INVTXWORD INVCRC
-------------------------------------------------------
u2d1  0        0        0       0        1         0
u2d2  0        0        0       0        30        0
u2d3  0        0        0       0        1         0
u2d4  0        0        0       0        30        0
u2d5  0        0        0       0        30        0
u2d6  0        6        0       0        30        0
u2d7  0        0        0       0        30        0
u2d8  0        0        0       0        30        0
u2d9  0        0        0       0        1         0

pass
```

Diagnosing an FC-AL Loop

This section describes how to diagnose an FC-AL loop problem. This section contains the following sub-sections:

FC-AL Loop Problem Indicators

The following symptoms indicate possible FC-AL loop problems:

1. The first indication observed by a customer might be performance degradation in the suspect array. See “Checking Performance Against Baseline Data” on page 107 for more detail.

2. A second indication might be Storage Automated Diagnostic Environment (StorADE) message monitoring from a host that is receiving remote array syslog messages. Storage Automated Diagnostic Environment monitoring can be configured to look for particular message classes in the log file that the array entries are written to. The program looks through this log file at a customer-determined frequency for the specified type of messages, and sends e-mail if a match is made. Typically, Storage Automated Diagnostic Environment message monitoring is configured to scan for warning or error messages. These message can also be examined in the array’s local syslog. The e-mail recipient can be the customer or any other destination the customer desires. See “Storage Automated Diagnostic Environment Message Monitoring” on page 108 for more detail.

3. A third indication of a problem may be a message or change of status in the Component Manager maintenance program GUI display, for example, a suspect FRU highlighted in red. Component Manager also sends e-mail to whomever the customer specifies and logs the failure into a customer-designated log file on the host that Component Manager is running on. See “Example syslog Error Messages” on page 109 for more details.

4. A fourth indication of a problem may be a warning or error log entry in the local array syslog file. Examine this file by using CLI commands via a Telnet or Tip connection. This file can also be transferred via ftp to another host for examination and archiving. See “Manual Examination of the syslog File” on page 108 and “Example syslog Error Messages” on page 109 for more details.

5. Additional indications of an FC-AL loop problem can provided by running the CLI commands described in “Normal Status” on page 96. See “Using CLI Diagnostic Commands” on page 110 for more detail.

If after this information has been gathered and examined and it has been determined that one of the back-end FC-AL loops has failed, but no definitive FRU an be identified, perform one or more of the diagnostic procedures described in the following sections.

Checking Performance Against Baseline Data

If the customer regularly runs a performance monitoring program where thresholds have been set, the `iostat` command shows whether one path to a Sun StorEdge T3+ array partner group is not performing to the established base line. For example:

CODE EXAMPLE 8-8 `iostat` Output for Normal (Baseline) Operation

```
r/s w/s Mr/s Mw/s wait actv wsvc_t asvc_t %w %b device
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c1t6d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t0d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t2d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c2t7d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c2t6d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c3t1d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c4t2d1
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c5t1d1
54.0 28.5 0.4 7.0 0.0 0.7 0.0 8.3 0 60 c5t1d0 (normal u1ctr I/O)
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c3t1d1
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c4t2d0
53.3 26.1 2.2 6.4 0.0 1.6 0.0 19.7 0 59 c6t2d1 (normal u2ctr I/O)
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c6t2d0
```

CODE EXAMPLE 8-9 `iostat` Output for Abnormal (Problem) Operation

```
r/s w/s Mr/s Mw/s wait actv wsvc_t asvc_t %w %b device
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c1t6d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t0d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t2d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c2t7d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c2t6d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c3t1d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c4t2d1
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c5t1d1
13.6 2.1 0.1 0.4 0.0 0.1 0.0 4.4 0 5 c5t1d0 (abnormal u1ctr I/O)
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c3t1d1
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c4t2d0
53.0 18.6 3.9 4.4 0.0 2.5 0.0 34.6 0 37 c6t2d1 (normal u2ctr I/O)
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c6t2d0
```

In the above example, if the normal `iostat` is used as a notification threshold, the impacted `iostat` indicates that there might be a problem in the master u1ctr controller in this redundant partner group.

Storage Automated Diagnostic Environment Message Monitoring

If Storage Automated Diagnostic Environment message monitoring is installed and running, it sends email messages indicating problems. For example, in the case of the performance impact illustrated above, the email might have the following data in it:

CODE EXAMPLE 8-10 Example Storage Automated Diagnostic Environment Message Monitoring Email Message Data

```
Mar 07 18:33:22 T3a ISR1[1]: W: u1d9 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 T3a ISR1[1]: W: u1d8 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 T3a ISR1[1]: W: u1d7 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 T3a ISR1[1]: W: u1d6 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 T3a ISR1[1]: W: u1d5 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 T3a ISR1[1]: W: u1d4 SVD_PATH_FAILOVER: path_id = 0
```

In this example, this data was pulled by Storage Automated Diagnostic Environment message monitoring from the remote host log file that the array sent `syslog` entries to. Storage Automated Diagnostic Environment message monitoring was run on that host and scanned the log looking for array log file messages of a `warning` or `error` class. The data in the example above indicates that drives u1d4-9 in the u1ctr controller completed a path failover from loop 1 (path 0) to loop 2 (path 1). This means that a hard failure or a threshold count was exceeded on the u1l1 loop. At this time, drives u1d4-9 are being serviced by the u1ctr only through the u1l2 loop. This is a good indication that there has been some kind of failure in the u1l1 interconnect card, the u1ctr controller, or one of the u1d1-9 drives.

Manual Examination of the `syslog` File

If Storage Automated Diagnostic Environment message monitoring is not running, the Sun StorEdge T3+ array CLI interface can be used to examine the unit's `syslog`. Use either the `cat` or `more` command on the log file. Either command outputs the complete log to the Telnet or Tip session screen. Alternatively, you can `ftp` the `syslog` file to the `telnet` or `tip` host and examine it with a text editor capable of

performing text searches with a character match. In the case of the example shown above, a search would be done for the error message type field of a `W:`. Such a search might display data similar to the following:

```
Mar 07 18:33:22 ISR1[1]: W: u1d9 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 ISR1[1]: W: u1d8 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 ISR1[1]: W: u1d7 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 ISR1[1]: W: u1d6 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 ISR1[1]: W: u1d5 SVD_PATH_FAILOVER: path_id = 0
Mar 07 18:33:22 ISR1[1]: W: u1d4 SVD_PATH_FAILOVER: path_id = 0
```

As can be seen, the data is the similar to what Storage Automated Diagnostic Environment message monitoring would display, and indicates the same possible failure condition on loop 1 (path 0).

Example `syslog` Error Messages

CODE EXAMPLE 8-11 displays some example `syslog` error messages that might indicate a back-end FC-AL drive loop problem.

CODE EXAMPLE 8-11 Drive Loop Problem Example Error Messages

```
Sep 27 18:36:53 T3A ROOT[1]:W:u1ctr Hardware Reset (1000) occurred
Sep 27 18:48:46 T3A ISR1[1]:W:SCSI Disk Error Occurred (path = 0x1, port =
0x1, lun = 0x0)
Sep 28 06:52:23 T3A CFGT[1]:W:u2ctr:Disabled
Sep 28 06:53:49 T3A LPCT[1]:E:u2ctr:Not present
Sep 28 06:53:49 T3A TMRT[1]:E:u2ctr:Missing; system shutting down in 30
minutes
Sep 28 06:53:49 T3A LPCT[1]:E:u2ctr:Not present
Sep 28 07:01:41 T3A ISR1[2]:W:u2d1 SCSI Disk Error Occurred (path = 0x1)
Sep 28 07:01:41 T3A ISR1[2]:W:Sense Data Description = Logical Unit Not
Ready, Initializing CMD Required
Sep 28 07:01:41 T3A ISR1[2]:W:u2d1 SCSI Disk Error Occurred (path = 0x1)
Sep 28 07:01:41 T3A ISR1[2]:W:Sense Key = 0x2, Asc = 0x4, Ascq = 0x2
Sep 28 07:01:41 T3A WXFT[2]:W:u2d1:Failed
Sep 28 07:01:41 T3A WXFT[2]:W:u2d1 hard err in vol (vol001) starting auto
disable
Sep 28 07:10:27 T3A LT01[1]:W:u2d1 Recon attempt failed
Sep 28 07:15:05 T3A ISR1[1]:W:SCSI Disk Error Occurred (path = 0x1, port =
0x1, lun = 0x0)
Sep 28 07:15:05 T3A ISR1[1]:W:Sense Data Description = SCSI Parity Error
Sep 28 07:18:03 T3A ISR1[1]:W:Sense Data Description = SCSI Parity Error
```

Using CLI Diagnostic Commands

Once the `syslog` file has been examined for warning or error messages and a conclusion is reached on which loop might have failed, other CLI commands can be used to verify or support that conclusion. These commands display the various status and current configuration of the loops.

Use a serial cable and Tip session to collect and analyze both controller's loop status information. The serial cable is necessary to see the loop configuration for the alternate controller, as the Telnet session only displays the current loop status as seen from the master controller.

For the example problem above, the CLI commands produce these results.

- The `fru stat` command would show a normal status for this failure.
- The `vol mode` command would show a normal status for this failure.
- The `port listmap` command would show a normal status for this failure.
- The `.loop stat` command would show a normal status for this failure.
- The `.disk pathstat` command would show a normal status for this failure.
- The `.disk linkstat` command would show the following error conditions for this failure.

CODE EXAMPLE 8-12 Example `.disk linkstat` Error Data

```
.disk linkstat u1d1-9 path 0 (master controller)
DISK LINKFAIL LOSSSYNC LOSSSIG PROTOERR INVTXWORD INVCRC
--------------------------------------------------------
u1d1 Disk Link Status Failed
u1d2 Disk Link Status Failed
u1d3 Disk Link Status Failed
u1d4 Disk Link Status Failed
u1d5 Disk Link Status Failed
u1d6 Disk Link Status Failed
u1d7 Disk Link Status Failed
u1d8 Disk Link Status Failed
u1d9 Disk Link Status Failed
fail
```

When the `.disk linkstat` command is run from the master controller, it is unable to access any of the link registers for drives u1d1-9. This supports the conclusion that loop 1 (path 0) has had a failure.

Once a suspect loop has been determined, use a process of elimination to locate the failed FRU on that loop as described in the following sections.

Using the ofdg Diagnostic Utility

If the problem is still unresolved, the last diagnostic tool to use is the off-line drive diagnostic utility (ofdg). Because the ofdg diagnostic requires that the T3+ partner group be removed from host access it is a highly disruptive procedure stops all data access to the T3+ partner group. Coordinate and schedule this down time with the customer.

To view the available ofdg utility command parameters, simply enter ofdg on the command line with no options.

```
:/:<15> ofdg
usage:  ofdg [-y] health_check
        ofdg [-y] fast_test u<encid>l{1|2}
        ofdg [-y] fast_find u<encid>l{1|2}
        ofdg [-y] find u<encid>l{1|2}
```

The ofdg parameters are:

- health_check does a fast Go / No-Go test of both loops using the current loop configuration. health_check uses fast_test, but no other parameters are required. (See “The health_check Option” on page 113 for additional details.)
- fast_test does a fast Go / No-Go test of the selected enclosure and loop with the current loop configuration. (See “The ofdg fast_test Option” on page 113 for additional details.)
- fast_find does a fast Go / No-Go test of the selected enclosure and loop. It also runs a simplified loop fault isolation diagnostic. (See “The ofdg fast_find Option” on page 114 for additional details.)
- find does an extensive Go / No-Go test. If loop failures are detected, it automatically initiates the full loop fault isolation diagnostic. This is similar to ondg find. (See “The ofdg find Option” on page 114 for additional details.)

See “Off-Line Drive Diagnostics and Replacement” on page 117 for a step-by-step description of using this utility to diagnose and replace a bad drive.

Caution – There are limitations to using the ofdg utility. Make sure you are aware of these limitations before running ofdg.

The following are limitations for using ofdg:

- Before running the ofdg utility, all disks other than those located in the u1 tray must be assigned to a LUN. Problems might occur if ofdg is run on systems where non u1 disks have not been assigned to volumes.
- ofdg does not detect missing loop cables.

- `ofdg` output goes to the `syslog` and serial port only.
- `ofdg` assumes at least one back-end loop cable is functional.

After installing a new drive, wait two minutes before running `ofdg`. Follow these steps to run `ofdg`:

1. **Perform an** `ofdg health_check` **operation.**

```
:/:<1> ofdg health_check
```

All loops are given either a Go or No-Go status.

- If there is a Go status, this indicates that the `ofdg` test did not detect any problems with the configuration and that there is no need for further tests.
- If there is a No-Go status, proceed to the next step.

2. **Perform an** `ofdg fast_test` **operation.**

```
:/:<2> ofdg fast_test u1l1
```

All loops are given either a Go or No-Go status.

- If there is a Go status, this indicates that the `ofdg` test did not detect any problems with the configuration and that there is no need for further tests.
- If there is a No-Go status, proceed to the next step.

3. **Perform an** `ofdg fast_find` **operation.**

```
:/:<3> ofdg fast_find u1l1
```

The loop is given a Go or No-Go status with progress indications. If a failure is reported on the first or nearest enclosure then the loop card in that enclosure should be swapped before repeating the test with the next unit.

If a failure is reported for the second (or further) enclosure, `fast_find` isolates the bad FRU(s) to either a bad interconnect cable or the two interconnect cards (which are connected to the interconnect cable in question). In this case, `fast_find` should be run from the partner controller to eliminate some FRUs.

If, after running `fast_find` in both directions, the problem has not been isolated to a single bad FRU, the bad FRU might be either the interconnect cable, the interconnect card, or both.

a. Replace the interconnect cable and retest.

b. Replace the interconnect card and retest

If the problem persists, continue to the next step.

4. **Perform an** `ofdg find` **operation.**

```
:/:<4> ofdg find u1l1
```

The loop is given a Go or No-Go status with progress indications. If a failure is detected, then Loop Fault Diag is automatically invoked to find the bad disk ports.

If `ofdg find` is not successful in solving the problem, the backplane should be suspected. See "Replacing the Chassis/Backplane Assembly" on page 126" for details.

The `health_check` Option

The `health_check` option provides a fast Go/No-Go Loop test for all the loops in the array. The `health_check` option calls `fast_test` multiple times, one time for each loop.

The `ofdg fast_test` Option

The `fast_test` option provides a fast Go/No-Go Loop test.

The `fast_test` option performs the following steps:

1. LAC_Reserve the FC-AL Loop device under test (DUT).
2. Test next nearest enclosure on Loop DUT.
3. Repeat Step 2 until all enclosures are tested.
4. LAC_Release the FC-AL Loop device under test (DUT).

The `fast_test` option uses only the two worst case data patterns as shown below:

```
#define ONDG_PATTERN_FOUR            0x7E7E7E7E     /* from SUN */
#define ONDG_PATTERN_SIX             0x4A4A4A4A     /* from SUN */
```

For each data pattern, the `fast_test` option performs the following:

- 2 synchronous Write/Read/Compares at 64-KB.
- 250 asynchronous Read/Writes at 64-KB.
- Monitors for errors (using all the FC-AL port counters on the Loop DUT, plus the counters from the single disk DUT).

The `ofdg fast_find` Option

The `fast_find` option provides a fast Go / No-Go Loop test (identical to `fast_test`), plus a simplified Loop Fault Diag.

The `fast_find` option performs the following steps:

1. LAC_Reserve the FC-AL Loop device under test (DUT)
2. Reconfigure Loop (via MUX) with next nearest enclosure on Loop DUT
3. Test next nearest enclosure on Loop DUT
4. Repeat Step 2 and Step 3 until all enclosures are tested

5) LAC_Release the FC-AL Loop device under test (DUT)

The big difference between `fast_find` and `find` is that `fast_find` does not attempt to drill down to a disk port (that is, detect and isolate down to a bad disk port), while `find` will try using Type 1 and Type 2 algorithms.

The `fast_find` option assumes that the probability of loop failures caused by either a bad interconnect cable or loop card is much higher than the probability of loop failures caused by a bad disk port. Therefore, `fast_find` should be used before `find` to first weed out bad interconnect cables and loop cards (then find should be used to weed out bad disk ports if problems still exist).

The `ofdg find` Option

The `find` option provides a Go / No-Go Loop test. If the loop test fails Loop Fault Diag is invoked to drill-down and find the bad FRU(s).

The `find` option uses two different Drill-down algorithms in order to detect bad FRU(s).

- Type 1—bypass one disk port at a time and test
- Type 2—find any three disk ports that work, then enable one disk port at a time and test. Use Type 2 only if Type 1 is unsuccessful.

The Loop Fault Diag has the capability to detect and isolate down to a single disk port but, depending on the system configuration, can be time consuming.

Repair Procedures

Begin by replacing the FRU that have the minimum impact to the customer's operation, as shown in the following order:

1. "Interconnect Card Replacement Procedure" on page 115
2. "RAID Controller Replacement Procedure" on page 116
3. "Off-Line Drive Diagnostics and Replacement" on page 117
4. "Chassis Replacement Procedure" on page 123

Interconnect Card Replacement Procedure

A single interconnect card can be removed without affecting the customer operation (assuming that the other card is working, of course). Data accessibility is maintained during the replacement and testing of a single interconnect card with no change in the host configuration.

For the example of a suspected loop 1 (path 0) problem, perform the following steps.

1. **From the CLI, disable the u1l1 interconnect card.**

```
:/:<1> disable u1l1
```

2. **When the u1l1 LED is flashing amber, remove and replace the interconnect card from the u1l1 position.**

   See "Removing and Replacing an Interconnect Card" on page 77

3. **From the CLI, enable the u1l1 interconnect card.**

```
:/:<2> enable u1l1
```

4. **Verify the repair by using the listed CLI status commands.**

   See "Using CLI Diagnostic Commands" on page 110.

5. **If this did not correct the problem, proceed to replacing the RAID controller as described in the next section.**

RAID Controller Replacement Procedure

If replacing the interconnect cards and cables did not resolve the loop 1 (path 0) problem, the next least-disruptive repair action is the removal and replacement of a RAID controller.

In a partner group, a single RAID controller card can be removed without denying access to all data (assuming appropriate multipathing software has been configured on the host). While data accessibility is maintained during the replacement and testing of a single RAID controller, performance is reduced during this procedure. The customer might elect to schedule the repair action during a time of reduced operations to the Sun StorEdge T3+ array system.

For the above example of a suspected loop 1 (path 0) problem, perform the following steps.

1. **From the CLI, disable the u1 RAID controller card.**

```
:/:<1> disable u1
```

   This causes a controller failover to the other controller. The Telnet session fails and the alternate controller becomes the master. VERITAS, if used, redirects the host I/O through the remaining path for the failed controller's volumes.

2. **When the u1 LED flashes amber, remove and replace the u1 controller card.**

   See "Removing and Replacing a Controller Card" on page 49.

3. **After the controller boots, verify the LED on u1 interconnect card is a solid green.**

4. **Restart a Telnet session to the array.**

5. **It may be necessary to disable and then enable the controller with the CLI commands to return it to service.**

   For example:

```
:/:<2> disable u1
:/:<3> enable u1
```

6. **Verify that VERITAS, if used, completes a path fail back to the replaced controller.**

   Consult your VERITAS documentation for VERITAS diagnostic procedures.

7. **Verify the repair by using the listed CLI status commands.**

   See "Using CLI Diagnostic Commands" on page 110.

8. **If replacing the u1 controller card does not correct the problem, replace the u2 RAID controller in the u2 enclosure.**

9. **If replacing the two RAID controllers does not correct the problem, proceed to replacing disk drives as described in the next section.**

Off-Line Drive Diagnostics and Replacement

If replacing the interconnect and RAID controller cards does not resolve the loop 1 (path 0) problem, the next step is to test, and if necessary replace, any suspect disk drives.

The test to use is the `ofdg` off-line diagnostic utility. The `ofdg` diagnostic requires the array partner group to be removed from host access. This is a highly disruptive procedure that stops all data access to the array. Coordinate and schedule this down time with the customer.

To administer and monitor the test, connect a serial maintenance cable and open a Tip session to the Sun StorEdge T3+ array.

The following steps describe how to test for the above example of a suspected loop 1 (path 0) problem.

Note – Before running the `ofdg` utility, all disks other than those located in the u1 tray must be assigned to a LUN. Problems may occur if `ofdg` is run on systems where non `u1` disks have not been assigned to volumes.

---

1. **Make sure that all disks other than u1 are assigned to a LUN.**

2. **Quiesce all I/O going to all volume(s) in that disk array and associated partner group.**

   Notify all applications to stop accessing any affected volumes. This may require stopping the application.

   Verify that all drive activity has stopped. The solid green drive activity LEDs indicate that the drives are idle.

3. **If the disk array is using any volume manager software, such as VERITAS, disable transactions to the volumes that reside on the array backplane you wish to replace and all other volumes in that partner group.**

   Consult the appropriate volume manager documentation for information on disabling the data hosts access to the array volumes.

4. **Unmount the volume(s) from the Solaris host.**

```
# unmount /T3-filesystem-name
```

5. **Unmount the internal array volume(s).**

```
:/:<4> vol unmount vol1
```

6. **Disconnect the fiber optic cables from the array controllers.**

7. **Establish a serial connection and Tip session to the** Master RAID controller **of the problem array.**

See “Establishing a Serial Port Connection” on page 7.

8. **Execute the** `set` **command and note the current values of** `logto` **and** `loglevel`**.**

```
:/:<1> set
bootmode  auto
bootdelay 3
sn        112035
ip        10.4.35.134
netmask   255.255.255.0
gateway   10.4.35.1
tftphost  123.123.123.6
tftpfile  releases/nb210/nb210p20.bin
hostname  qatest
timezone  GMT+00
vendor    0301
model     501-5710-02(51)
revision  020100
logto     /Aug9
loglevel  3
rarp      off
mac       00:20:f2:00:03:b9
```

You will need these values in step Step 18.

9. **From the host Tip session, set the** `logto` **to** `1` **and the** `loglevel` **to** `4`**.**

```
:/:<5> set logto 1
:/:<6> set loglevel 4
```

These settings display all messages to the Tip session screen. The output includes all messages from `information` up to `error`.

10. **Run a** `find` **test against loop 1.**

```
:/:<7> ofdg find u1l1
WARNING - Volume data will be offline while OFDG is running.
Continue ? [N]: y
```

How far the test has go into the loop to identify the failed FRU determines how long the test runs. The `find` test may also have to be run again with the `u2l1` parameter if no failures are found with the `u1l1` parameter.

11. **Examine the output in detail to identify the failed FRU.**

For comparison, a test run that found no errors is shown in CODE EXAMPLE 8-13. (This test might take 8 minutes to complete.)

CODE EXAMPLE 8-13 `ofdg` Sample Output (No Errors)n

```
:/:<8> ofdg find u1l1
WARNING - Volume data will be offline while OFDG is running.
Continue ? [N]: y
ONDG Initiated
FIND Initiated on u1l1
Loop 1 Configured as  <1>
Loop 2 Not Available
Loop 1 Configured as  <1+>
Loop 2 Not Available
Loop 1 Configured as  <1>
Loop 2 Not Available
Loop 1 Configured as  <1+>
Loop 2 Not Available
FIND Completed on u1l1
STATUS = PASS
u1 PASS
ONDG Completed
```

```
In Syslog:
May 26 19:18:03 pshc[1]: N: ofdg find u1lp1
May 26 19:18:18 pshc[1]: N: ofdg find u1l1
May 26 19:18:22 ONDG[1]: N: ONDG Initiated
May 26 19:18:22 ISR1[1]: N: u1ctr ISP2200[2] Received LIP(f7,e8) async event
May 26 19:18:22 FCC0[1]: N: u1ctr Port event received on port 0, abort 0
May 26 19:18:23 BELP[1]: N: u1l1 ONDG Loop Fault Diag Initiated
May 26 19:18:28 CFGT[1]: N: u1ctr: Reserved A Loop: A Mask=<1>; B Mask=<1>
May 26 19:18:38 ISR1[1]: W: u1d4 SVD_PATH_FAILOVER: path_id = 0
May 26 19:18:38 ISR1[1]: W: u1d5 SVD_PATH_FAILOVER: path_id = 0
May 26 19:18:38 ISR1[1]: W: u1d6 SVD_PATH_FAILOVER: path_id = 0
May 26 19:18:38 ISR1[1]: W: u1d7 SVD_PATH_FAILOVER: path_id = 0
May 26 19:18:38 ISR1[1]: W: u1d8 SVD_PATH_FAILOVER: path_id = 0
May 26 19:18:38 ISR1[1]: W: u1d9 SVD_PATH_FAILOVER: path_id = 0
May 26 19:18:41 LPCT[1]: N: u1ctr: ISP not ready on loop 1
May 26 19:18:50 LPCT[1]: N: u1d1: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d2: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d3: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d4: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d5: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d6: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d7: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d8: Bypassed on loop 1
May 26 19:18:50 LPCT[1]: N: u1d9: Bypassed on loop 1
May 26 19:19:11 LPCT[1]: N: u1ctr: ISP not ready on loop 1
May 26 19:19:31 ISR1[1]: N: u1ctr ISP2200[0] Received LIP(f8,d1) async event
May 26 19:19:32 LPCT[1]: N: u1d1: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d2: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d3: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d4: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d5: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d6: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d7: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d8: Not bypassed on loop 1
May 26 19:19:32 LPCT[1]: N: u1d9: Not bypassed on loop 1
May 26 19:19:33 SVDT[1]: N: 9 fcal ports were detected on l1
May 26 19:19:41 ISR1[1]: N: u1ctr ISP2200[0] Received LIP(f7,ef) async event
May 26 19:19:43 SVDT[1]: N: 10 fcal ports were detected on l1
May 26 19:20:04 BELP[1]: N: u1l1 ONDG No Loop Trouble Found
```

```
May 26 19:20:04 CFGT[1]: N: u1ctr: Release A Loop: A Mask=<1>; B Mask=<1>
May 26 19:20:05 ISR1[1]: N: u1d4 SVD_PATH_FAILBACK: path_id = 1
May 26 19:20:05 ISR1[1]: N: u1d5 SVD_PATH_FAILBACK: path_id = 1
May 26 19:20:05 ISR1[1]: N: u1d6 SVD_PATH_FAILBACK: path_id = 1
May 26 19:20:05 ISR1[1]: N: u1d7 SVD_PATH_FAILBACK: path_id = 1
May 26 19:20:05 ISR1[1]: N: u1d8 SVD_PATH_FAILBACK: path_id = 1
May 26 19:20:05 ISR1[1]: N: u1d9 SVD_PATH_FAILBACK: path_id = 1
May 26 19:20:10 BELP[1]: N: u1l1 ONDG Loop Fault Diag Completed
May 26 19:20:10 ONDG[1]: N: FIND Initiated on u1l1
May 26 19:20:10 ONDG[1]: N: FIND Completed on u1l1
May 26 19:20:10 ONDG[1]: N: STATUS = PASS
May 26 19:20:10 ONDG[1]: N: u1 PASS
May 26 19:20:13 ONDG[1]: N: ONDG Completed
```

CODE EXAMPLE 8-14

12. **Once the failed disk drive FRU has been identified, remove the suspect disk drive from the configuration with the** `vol disable` **command.**

```
:/:<9> fru s u1d1-9

DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      30    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      29    vol1
u1d6    ready    enabled     data disk   ready      ready      29    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      32    vol3

:/:<10> vol disable u1d9

:/:<11> fru s u1d1-9
DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      30    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      29    vol1
u1d6    ready    enabled     data disk   ready      ready      29    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      32    vol3
```

13. **Remove and replace the suspect disk drive from the enclosure.**

    See “Repairing Disk Drives” on page 64.

    The drive spins up and the `sysarea` data copies to it from another drive in the u1 enclosure. After the copy is complete, a volume reconstruction starts.

14. **Rerun the** `ofdg find` **diagnostic through the suspect loop as described in Step 10 and Step 11.**

    Once the test completes, examine and compare the two outputs to insure that the fault has been corrected.

    a. **If the problem is resolved, proceed with Step 16 through Step 15.**

b. If the problem is not resolved, proceed with Step 16 through Step 15 and then replace the backplane chassis.

See "Chassis Replacement Procedure" on page 123 and "Replacing the Chassis/Backplane Assembly" on page 126".

15. Remount the volumes.

```
:/:<14> vol mount vol1

:/:<15> vol stat

vol1          u1d1   u1d2   u1d3   u1d4   u1d5
mounted        0      0      0      0      0
```

16. Restart the volume reconstruction with the `vol recon` command on the replaced disk drive.

```
:/:<16> vol recon u1d9
```

17. Reconnect the fibre optic cable to the MIAs,

18. Reset `logto` and `loglevel` to the original values noted in Step 8.

```
:/:<17> set logto *
:/:<18> set loglevel 3
```

Chassis Replacement Procedure

If none of the above procedures resolve the problem, the next repair action is replacement of the chassis/backplane assembly. A replacement part must be on site before beginning this procedure. Before starting, the customer must off load all the data that is contained in the array. The array must then be removed from host operation.

The procedure for replacing a backplane is described in "Replacing the Chassis/Backplane Assembly" on page 126.

Once the backplane has been replaced and the previous FRUs installed, the `ofdg` diagnostics need to be rerun.

If the problem persists, replace the entire Sun StorEdge T3+ array.

CHAPTER 9

Chassis / Backplane Assembly

This chapter describes how to replace the chassis / backplane assembly and contains the following sections:

Troubleshooting the Chassis / Backplane Assembly

The array chassis FRU rarely needs to be replaced. However, the chassis part number is available to replace the backplane and chassis if necessary. (These must be replaced together, because they are factory aligned.)

Problems with the backplane most likely occur because of an electrical short or a bent or broken pin connector. These problems first appear as a failure of another FRU component, such as an interconnect failure or drive failure. If replacing the FRU that appears to be failed does not correct the problem, then examine the backplane connector that the FRU connects to for bent or broken pins. If nothing is obvious, then install another spare FRU to verify that a failed FRU component is not causing the problem. If all possibility of a FRU component failure has been eliminated and the problem still remains, it is likely to be a backplane failure.

Replacing the Chassis / Backplane Assembly

If there is a backplane failure, replace it with the following procedure.

Caution – Replacing a Sun StorEdge T3+ array chassis interrupts array operation.

Note – If the Sun StorEdge T3+ array is part of a partner group, access to all volumes in the partner group is unavailable during this backplane replacement procedure. Assess the impact of unmounting volumes and stopping applications prior to starting this procedure.

1. **Perform full backups of data on affected partner groups for all accessible volumes.**
2. **From the data hosts, quiesce all I/O going to all volume(s) in that disk array and associated partner group.**
   - Notify all applications to stop accessing any affected volumes by unmounting the volume(s) or stopping the application if necessary.
   - Verify that all drive activity has stopped. The drive activity LEDs become solid green, indicating that the drives are idle.
3. **If any volume manager software is running, such as VERITAS, disable transactions to the volumes that reside on the Sun StorEdge TT3+ array backplane you are replacing and to all other volumes in that partner group.**
   - Consult the appropriate volume manager documentation for information on disabling the data hosts access to the Sun StorEdge T3+ array volumes.
4. **Execute the** `shutdown` **command.**

```
:/:<1> shutdown
Shutdown the system, are you sure? [N]: y
```

5. **Power down the failed disk array.**

Press the power button once on each power and cooling unit to turn the switch off (FIGURE 9-1).

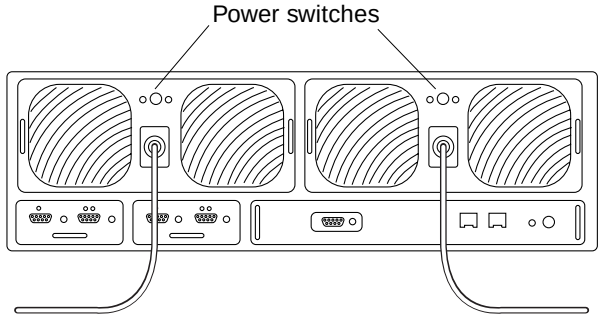

FIGURE 9-1 Power Switch Location

All arrays power down automatically when any one array in the partner group is powered down.

6. **Record the Sun StorEdge T3+ array system serial number and MAC address.**

Locate the pull-out tab at the left side of the array next to the first disk drive as shown in FIGURE 9-2. This tab contains the array serial number and media access control (MAC) address. The serial number is located on the top left portion of the pull-out tab and begins with the part number 595-xxxx. Record this information to transcribe it onto the new chassis.

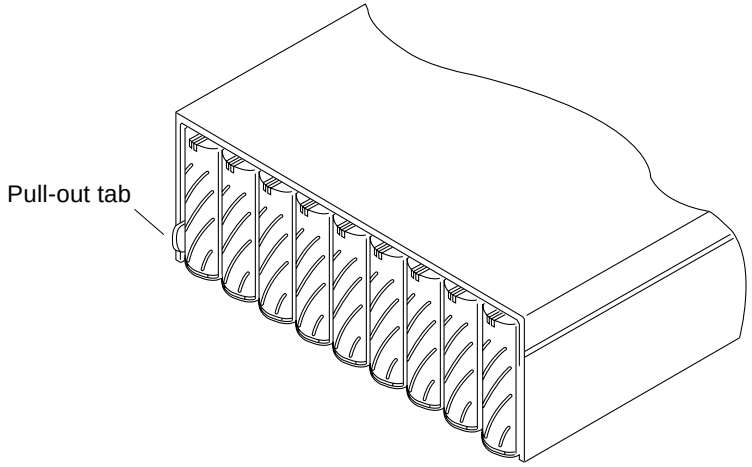

FIGURE 9-2 Serial Number and MAC Address on Pull-out Tab

7. On the Sun StorEdge T3+ array, disconnect all external cables.

Disconnect all power, interconnect, host FC-AL, MIA, and Ethernet cables.

Note – If the array is part of a partner group, note down the placement of the host FC-AL connections and loop cables. You need this information in Step 12.

8.Remove the chassis if it is mounted in a cabinet:

a. Remove the two screws at the back of the chassis that secure it to the side rails in the cabinet.

b. Slide the chassis out of the cabinet (FIGURE 9-3**).**

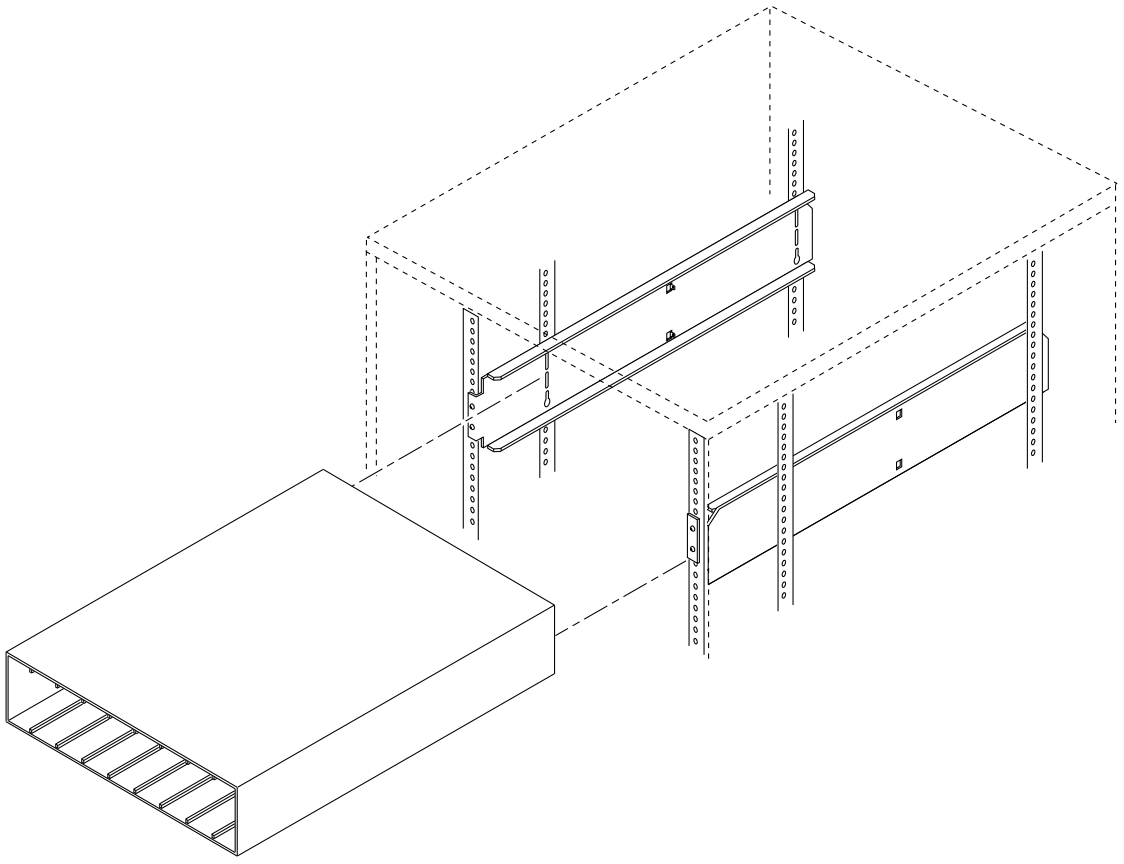

FIGURE 9-3 Removing the Chassis

9. **Move the failed array to an area that both the front and back can be easily accessed.**

Caution – Use two people to lift and move the array. It can weight up to 67 lbs (30 kg).

10. **One at a time, remove all the FRU components from the failed chassis and put them in the new chassis, ensuring same location placement.**

Caution – FRUs are extremely sensitive to static electricity. Use proper antistatic wrist strap and procedures when handling any FRU. Observe all static electricity precautions as described in “Static Electricity Precautions” on page 5.

a. **Unlock a FRU by pushing in on the latch handle(s) with a coin or small screwdriver to release the latch handle(s).**

b. **Pull the FRU straight out.**

c. **Lock the FRU back into place in the new chassis by pushing in and securing the latch handle(s) with a coin or small screwdriver.**

Caution – Maintain disk positions or data could be lost.

d. **Remove and replace the controller card.**

See “Removing and Replacing a Controller Card” on page 49 for instructions.

e. **Remove and replace the interconnect cards.**

See “Removing and Replacing an Interconnect Card” on page 77 for instructions.

f. **Remove and replace the power and cooling units.**

See “Removing and Replacing a Power and Cooling Unit” on page 85 for instructions.

g. **Remove and replace the disk drives.**

See “Repairing Disk Drives” on page 64 for instructions.

Note – When removing disk drives, label each one with its slot position in the unit so you can be replace the drives in the correct slots.

11. **Move the replacement chassis back into place.**

- If you are mounting the chassis in a cabinet:

- Prepare for the new chassis by installing the base plate. Use the base plate from the old chassis.
- Align the new chassis with the side rails and slide the chassis into the cabinet.
- Replace the two screws at the back of the chassis to secure the chassis to the cabinet.

12. **Connect all cables previously removed, but** *do not* **power up the array(s).**

Note – If the array is part of a partner group, make sure that the host FC-AL cables are recabled to the same Sun StorEdge T3+ array FC-AL connections that they were removed from as you noted down in Step 7. Also ensure that loop cables are properly recabled.

---

13. **Add the T3+ array serial number and the MAC address to the new chassis.**

    Locate the pull-out tab at the left side of the array next to the first disk drive. Use a fine-tipped permanent marker to write the information on this tab—you also need the information for the next two steps.

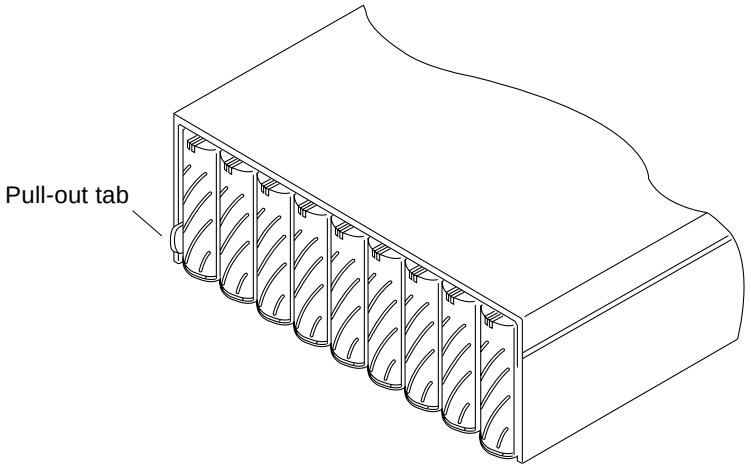

FIGURE 9-4 Serial Number and MAC Address on Pull-out Tab

14. **Contact the appropriate Contract Administrator (CA) of the Contracts Verification Group (CVG) to relay the system serial number and new chassis information.**

15. **On the RARP server, update the `/etc/ethers` file.**

    Replace the MAC address entry of the failed chassis with the MAC address of the new chassis. For example:

    ```
    8:0:20:6d:93:7e array-name
    ```

    In this example:

    - `8:0:20:6d:93:7e` is the new MAC address
    - *array-name* is the name of the old array.

    Note that if the failed unit was an alternate master, the unit's MAC address may not be in the `/etc/ethers` file. In this case, no file changes are required.

16. **Verify that the `/etc/hosts` file contains the previous IP address and array name.**

    For example:

    ```
    192.129.122.111 array-name
    ```

    In this example, `192.129.122.111` is the IP address assigned previously.

17. **Verify that the `/etc/nsswitch.conf` file on the RARP server references the local system files.**

    To ensure that the Solaris software environment uses the changes made to the `/etc/ethers` and `/etc/hosts` files, edit the `host` and `ethers` entries in the `/etc/nsswitch.conf` file so that the `files` parameter appears before the `[NOTFOUND=return]` statements. For example:

    ```
    hosts:      nis files [NOTFOUND=return]
    ethers:     nis files [NOTFOUND=return]
    ```

18. **Ensure that the RARP daemon is running on the RARP server:**

    ```
    rarpserver# ps -eaf | grep rarpd
    ```

19. **If the RARP daemon is not already running on the RARP server, start it by entering:**

    ```
    rarpserver# /usr/sbin/in.rarpd -a &
    ```

20. **Verify that AC power is present on each of the chassis' power and cooling units.**

    The AC LED on each power and cooling unit glows solid amber and the fans turn at low speed.

21. **Press the power button on the power and cooling units to power on the array(s).**

    FIGURE 9-1 shows the power button location. The AC and power supply (PS) LEDs on the power and cooling units show green.

    After you power on, the Sun StorEdge T3+ array JumpStart™ feature reassigns the array's previous IP address to the new MAC address.

    Allow time to complete the boot cycle. When all LEDs are green, proceed to the next step.

22. **Check the LEDs at the front and back of the unit to ensure that all components are receiving power and are functional.**

    While the drives are spinning up, the LEDs blink. The array boot time take up to several minutes, after which all LEDs should be solid green, indicating that the unit is receiving power and there is no drive activity.

Note – The batteries in the power and cooling units recharge after powering on the unit. While the batteries are recharging, write-behind cache is disabled.

Note – If the green power and cooling unit LEDs on connected units do not light, press the power switches on those units.

---

23. **Use the CLI to verify that all components are functioning properly.**

    To verify status using the CLI, open a Telnet session to the disk array and verify volume and FRU status. as described "Checking FRU Status" on page 35.

```
:/:<1> fru stat
:/:<2> vol stat
```

    Note that when the backplane is replaced, the data host volume's WWN changes. The WWN is derived from the backplane serial number. Since the volume WWN is part of the volume's device path on the data host, the device path definition on the data host changes. Therefore, you must reconfigure the data host to recognize the new WWNs.

24. **Configure the data host to recognize the new WWNs by executing the following command on the data host:**

```
datahost# drvconfig; disks; devlinks
```

Note – Any applications specifically dependent on the volume's device path also need to be changed. Refer to each application's documentation for instructions.

---

25. **Execute a `format` command on the data host to verify that the Sun StorEdge T3+ array devices are seen.**

    The Sun StorEdge T3+ array volumes are now usable by the data host and can be mounted or re-enabled with the appropriate volume manager software.

CHAPTER 10

Hardware Reconfiguration

This chapter provides procedures for reconfiguring existing array hardware to create new configurations. It includes the following sections:

Connecting Single Controller Units to Form a Partner Group

Caution – This procedure destroys data. Back up all your data before beginning this procedure.

This section describes how to reconfigure two existing single controller units that contain data to form a partner group (redundant controller units). You will need two interconnect cables to connect the units. See Appendix A for a part number and illustration of the interconnect cable.

This procedure includes the following sections:

Preparing the arrays

1. **Decide which unit is the master controller and which is the alternate master.**
2. **Back up the data on both arrays.**

Caution – Make sure you back up data on both units *before* proceeding! You need to re-create the volume(s) on the alternate master after cabling the units together.

3. **Ensure that the data path between the host and both arrays has been quiesced.**

   There must not be any I/O activity.

4. **Start a Telnet session with both arrays.**

   a. **On the host, use the** `telnet` **command with the array name (or IP address) to connect to the array.**

```
# telnet array_name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

Telnet session (129.150.47.101)
```

   b. **Log in to the array by typing** `root` **and your password at the prompts.**

      The array prompt is displayed.

5. **Verify that firmware levels for all array firmware are the same on the master unit and alternate master unit.**

   On both arrays:

   a. **Type** `ver` **to display the controller firmware level.**

      For example:

```
:/:<1> ver

T3B Release 2.1 2002/07/30 19:16:42 (10.4.35.134)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.
```

      In this example, the controller firmware level is listed as `Release 2.1`

b. Type `fru list` to display EPROM, disk drive, and interconnect card firmware levels.

For example:

```
:/:<2> fru list
ID      TYPE                VENDOR       MODEL          REVISION       SERIAL
------  ------------------  -----------  -----------  -------------  --------
u1ctr   controller card     0301         501-5710-02( 020100/020101  112035
u2ctr   controller card     0301         501-5710-02( 020100/020101  112122
u1d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMKJ
u1d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH2A
u1d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1H9WS
u1d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM64
u1d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMC2
u1d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM63
u1d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1HE3A
u1d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNK0
u1d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM5P
u2d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HHH5
u2d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMJC
u2d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1HGKR
u2d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HLBJ
u2d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNH0
u2d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH4Z
u2d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1H92W
u2d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HN9T
u2d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HKQP
u1l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    1413
u1l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    2294
u2l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    001415
u2l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    002054
u1pcu1  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001787
u1pcu2  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001784
u2pcu1  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001544
u2pcu2  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001545
u1mpn   mid plane           SCI-SJ       370-3990-01- 0000           000953
u2mpn   mid plane           SCI-SJ       370-3990-01- 0000           000958
```

In this example:

- EPROM firmware version is listed as `Controller card, Revision 020100/020101`
- Disk drive firmware version is listed as `Revision A726`
- Interconnect card (`loop card`) firmware version is listed as `Revision 5.02 Flash`

c. Upgrade firmware, if necessary.

- If the firmware levels are the same on each unit, then proceed to Step 6.
- If the firmware levels for any of the four types of firmware are different between the master and alternate master, upgrade the firmware that does not match on both units. Refer to the upgrading firmware instructions in the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

6. **On both units, use the `set -z` command to return critical array settings to factory defaults.**

   When prompted to respond, answer `y` (yes). For example:

```
:/:<3> set -z
WARNING - Resetting system NVRAM to default, are you sure? [N]: y
t300:/:<4>
```

Note – The `set -z` command resets the `set` parameters, where as `boot -w` wipes out all the volumes and `sys` parameters. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for more detailed information on setting block size.

Caution – The `set -z` command resets the IP address of the units to 0.0.0.0. You will to reassign the IP address to the master unit after you cable the partner group together, but *before* powering on, as described in the next section.

---

7. **Power off both units.**

   a. **Type:**

```
:/:<4> shutdown
Shutdown the system, are you sure? [N]: y
```

   b. **Press the power button once on each power and cooling unit to turn the switch off.**

Cabling a Partner Group

After changing the array settings on the alternate master to the factory default and reverifying that both units run the same firmware levels, you are ready to connect the arrays.

1. **Place the alternate master on top of the master unit.**

- If the units are installed in a cabinet, make sure that the alternate master is installed in the slot directly above the master unit. If you need to change the position in the cabinet, refer to the rackmount installation instructions in the *Sun StorEdge T3+ Array Installation and Configuration Manual*.
- If the units are cabled to the hosts and power sources such that they cannot be placed in close proximity, rearrange the cabling so that the units can be placed together.

2. **Make sure that the 100BASE-T cables are connected to a network with the same management host.**

3. **Connect the interconnect cables to the interconnect cards as shown in** FIGURE 10-1**.**

   Make sure you connect the cables to the correct interconnect card connectors *exactly* as shown in the figure. This cable connection determines the master and alternate master relationship. Tighten the retaining screws.

   The remaining connectors are reserved for possible future expansion units.

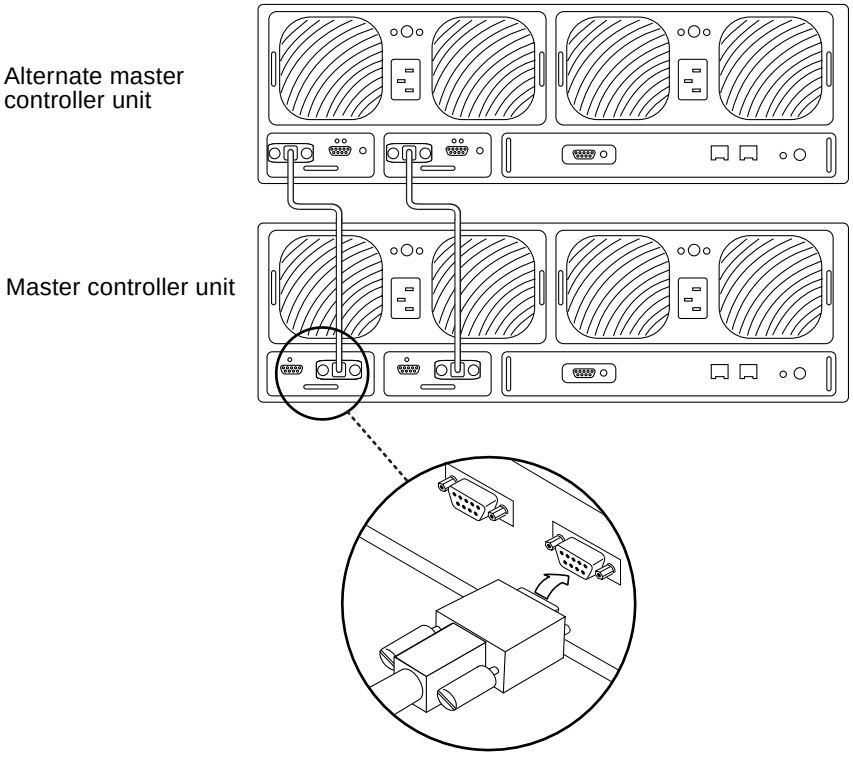

FIGURE 10-1 Connecting the Interconnect Cables

A fully cabled partner group is shown below.

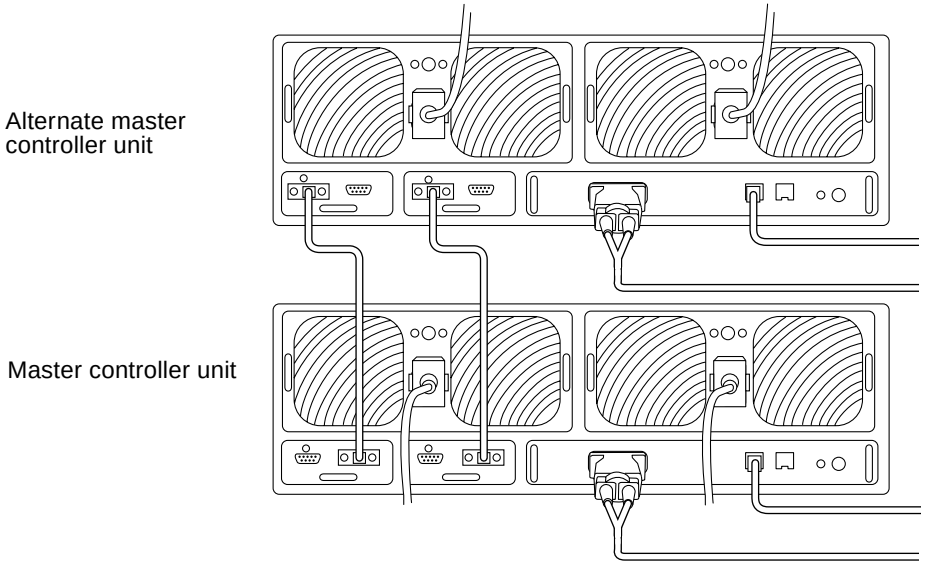

FIGURE 10-2 Fully Cabled Partner Group

Caution – Do not power on the arrays yet. You must configure a RARP server, connected to the array, with the IP address *before* powering on.

Establishing a New IP Address

The JumpStart feature automatically downloads a newly assigned IP address to the array. To enable this feature, you must edit your host file on a RARP server before powering on the array. After you power on, the IP address is automatically assigned. Before you begin, make sure you have the following:

- **MAC address**. The MAC address is located in the pullout tab at the front of the array (FIGURE 10-3).

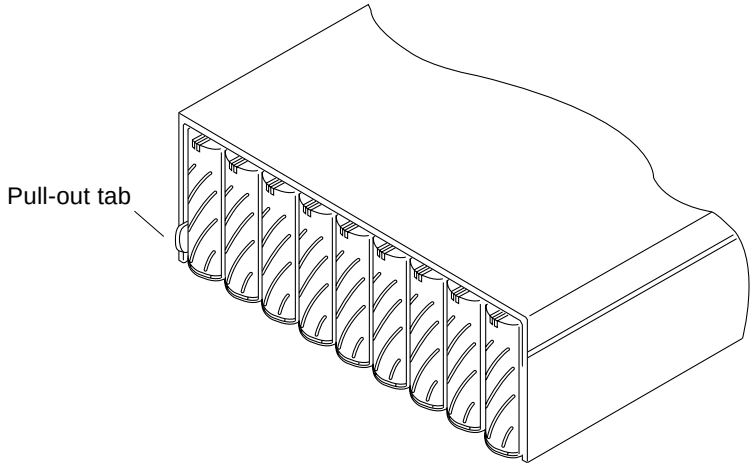

FIGURE 10-3 Location of Pull-Out Tab With MAC Address

- **IP address.** For this information, contact the person who maintains your network.
- **Array name.** This is the user-assigned name of the array.

To set the network IP address for the array:

1. **On a host connected to the same subnet as the array, edit the** `/etc/ethers` **file by adding the MAC address and array name.**

   For example:

   ```
   8:0:20:7d:93:7e array-name
   ```

   In this example:

   - `8:0:20:7d:93:7e` is the MAC address.
   - *array-name* is the name of the array you are installing.

2. **Edit the** `/etc/hosts` **file with the IP address and array name.**

   For example:

   ```
   192.129.122.111 array-name
   ```

   In this example, `192.129.122.111` is the assigned IP address.

3. **Edit the** `/etc/nsswitch.conf` **file to reference the local system files.**

   To ensure that the Solaris software environment uses the changes made to `/etc/ethers` and `/etc/hosts` files, edit the `host` and `ethers` entries in the `/etc/nsswitch.conf` file so that the `files` parameter appears before the `[NOTFOUND=return]` statements.

```
hosts:      nis files [NOTFOUND=return]
ethers:     nis files [NOTFOUND=return]
```

4. **Determine if the RARP daemon is running by typing:**

```
# ps -eaf | grep rarpd
```

   - If the RARP daemon is running, proceed to Step 6.
   - If the RARP daemon is not running, proceed to the next step.

5. **Start the RARP daemon in the Solaris software environment by typing:**

```
# /usr/sbin/in.rarpd -a &
```

6. **Power on both arrays by pressing the power button on each power and cooling unit.**

   All power and cooling unit LEDs on both units turn green, indicating that power has been restored. The IP address automatically downloads to the master controller unit after you power on.

Note – In some cases, the array times out before receiving the RARP request through an Ethernet switch. If this time out happens, the array cannot receive the assigned IP address. An improper spanning-tree setting of the Ethernet switch might cause this time out. Refer to your switch vendor documentation for information on spanning-tree settings and how to change them. Changing this setting properly enables the array to receive the RARP request before timing out.

Defining and Mounting Volumes on the Alternate Master

Once the units are cabled and power has been restored to both units successfully, define and mount the volume(s) on the alternate master.

Note – Make sure that both units are online and that all LEDs are green. It can take several minutes after powering on for the units to be ready.

---

1. **Start a Telnet session with the master controller unit.**

   a. **On the host, use the** `telnet` **command with the array name (or IP address) to connect to the master unit.**

```
# telnet disk_tray_name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

Telnet session (129.150.47.101)
```

Note – The Telnet session verifies that your network connection is good. If you cannot connect through the Telnet session, you might have miscabled the partner group. See “Identifying Miscabled Partner Groups” on page 36 to determine if this is the problem. If the partner group is cabled correctly, then the IP address might not be assigned correctly. If you suspect this as the problem, verify the IP address in a serial cable connection and verify that the RARP server is functional.

---

   b. **Log in to the array by typing** `root` **and your password at the prompts.**

      The array prompt is displayed.

2. **Check the FRU status using the `fru list` and `fru stat` commands.**

   Make sure that all FRUs are displayed and that FRU conditions are good, as shown in the following examples:

```
:/:<1> fru list
ID      TYPE                VENDOR       MODEL        REVISION       SERIAL
------  ------------------  -----------  -----------  -------------  --------
uID      TYPE                VENDOR       MODEL        REVISION       SERIAL
------  ------------------  -----------  -----------  -------------  --------
u1ctr   controller card     0301         501-5710-02( 020100/020101  112035
u2ctr   controller card     0301         501-5710-02( 020100/020101  112122
u1d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMKJ
u1d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH2A
u1d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1H9WS
u1d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM64
u1d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMC2
u1d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM63
u1d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1HE3A
u1d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNK0
u1d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM5P
u2d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HHH5
u2d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMJC
u2d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1HGKR
u2d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HLBJ
u2d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNH0
u2d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH4Z
u2d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1H92W
u2d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HN9T
u2d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HKQP
u1l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    1413
u1l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    2294
u2l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    001415
u2l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    002054
u1pcu1  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001787
u1pcu2  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001784
u2pcu1  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001544
u2pcu2  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001545
u1mpn   mid plane           SCI-SJ       370-3990-01- 0000           000953
u2mpn   mid plane           SCI-SJ       370-3990-01- 0000           000958
```

```
:/:<2> fru stat
CTLR    STATUS   STATE       ROLE        PARTNER    TEMP
------  -------  ----------  ----------  -------    ----
u1ctr   ready    enabled     master      u2ctr      31.0
u2ctr   ready    enabled     alt master  u1ctr      30.5

DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      30    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      29    vol1
u1d6    ready    enabled     data disk   ready      ready      29    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      32    vol3
u2d1    ready    enabled     data disk   ready      ready      34    vol2
u2d2    ready    enabled     data disk   ready      ready      38    vol2
u2d3    ready    enabled     data disk   ready      ready      36    vol2
u2d4    ready    enabled     data disk   ready      ready      37    vol2
u2d5    ready    enabled     data disk   ready      ready      34    vol2
u2d6    ready    enabled     data disk   ready      ready      36    vol4
u2d7    ready    enabled     data disk   ready      ready      35    vol4
u2d8    ready    enabled     data disk   ready      ready      40    vol4
u2d9    ready    enabled     data disk   ready      ready      36    vol4

LOOP    STATUS   STATE       MODE        CABLE1     CABLE2     TEMP
------  -------  ----------  -------     ---------  ---------  ----
u2l1    ready    enabled     master      installed  -          29.5
u2l2    ready    enabled     slave       installed  -          31.0
u1l1    ready    enabled     master      -          installed  29.5
u1l2    ready    enabled     slave       -          installed  30.0

POWER   STATUS   STATE       SOURCE  OUTPUT  BATTERY  TEMP    FAN1    FAN2
------  -------  ---------   ------  ------  -------  ------  ------  ------
u1pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u1pcu2  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu2  ready    enabled     line    normal  normal   normal  normal  normal
```

Note – The batteries in the power and cooling units recharge after powering on the unit. During the recharge, a `fault` message is displayed in the `fru stat` output for the batteries. While the batteries are recharging, write-behind cache is disabled.

3. **Use the** `vol add` **command to create the volume(s) on the alternate master as follows:**

   a. **Define the volume name (**`vol add` *volume-name***).**

   b. **Define the drives (**`data u2d`*n*`-`*n***) on which the volume resides, where:**

      - `u2` is the array unit number.
      - `d`*n*`-`*n* are the disk drives, *n* = 1 to 9.

   c. **Define the RAID level (**`raid` *n***), where** *n* **= 0, 1, or 5.**

   d. **Optional: define the hot spare drive (**`standby u`*n*`d9`**) where:**

      - `u2` is the array unit number.
      - `d9` is the number of the hot spare disk drive .

```
:/:<3> vol add volume-name data undn-n raid n standby und9
```

   For example:

```
:/:<4> vol add vol2 data u2d1-8 raid 5 standby u2d9
```

      - `vol2` is the volume name.
      - `u2d1-8` indicates the location of the volume: unit 2, drive 1 through 8
      - `raid 5` is RAID level 5.
      - `standby u2d9` is the location of the hot spare: unit 2, drive 9.

4. **Check the status of the volumes to ensure that you created the volume correctly.**

   The status of all drives must be `0`. For example:

```
:/:<1> vol stat
vol2          u2d1   u2d2   u2d3   u2d4   u2d5   u2d6   u2d7   u2d8   u2d9
unmounted      0      0      0      0      0      0      0      0      0

:/:<1> vol list
volume         capacity     raid  data      standby
vol1           236.058 GB   5     u1d1-8    u1d9
vol2           236.058 GB   5     u2d1-8    u2d9
```

5. **Initialize the volumes.**

```
:/:<3> vol init vol1 data
:/:<3> vol init vol2 data
```

6. **Mount the volumes.**

```
:/:<3> vol init vol1
:/:<3> vol init vol2
```

7. **Use the `format` command on a Solaris host to find out infromation about the new volume.**

   The `format` command probes for new devices and provides information about them including their sizes and pathnames. Refer to the `format` man page for more information on this ocmmand.

Disconnecting a Partner Group to Form Single Controller Units

Caution – Back up all data before beginning this procedure.

This section describes how to reconfigure a partner group to form two existing single controller units.

This procedure includes the following sections:

Preparing the Arrays

1. **Back up all data on the partner group.**

Caution – Make sure you back up data *before* proceeding.

2. **Ensure that the data path between the host and the partner group has been quiesced.**

   There must not be any I/O activity.

3. **Start a Telnet session with the master unit.**

   a. **On the host, use the `telnet` command with the array name (or IP address) to connect to the array.**

```
# telnet array_name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

Telnet session (129.150.47.101)
```

b. Log in to the array by typing `root` and your password at the prompts.

The array prompt is displayed.

4. View a listing and the status of the volumes.

```
:/:<1> vol list
volume          capacity     raid  data      standby
vol1            236.058 GB   5     u1d1-8    u1d9
vol2            236.058 GB   5     u2d1-8    u2d9

:/:<2> vol stat
volume          capacity     raid  data      standby
vol1            236.058 GB   5     u1d1-8    u1d9
vol2            236.058 GB   5     u2d1-8    u2d9
```

5. Unmount vol1.

```
:/:<1> vol unmount vol1
```

6. Remove vol1.

```
:/:<1> vol remove vol1
```

7. Power off both units.

a. Type:

```
:/:<5> shutdown
Shutdown the system, are you sure? [N]: y
```

b. Press the power button once on each power and cooling unit to turn the switch off.

8. Remove the interconnect cables from the back of each array.

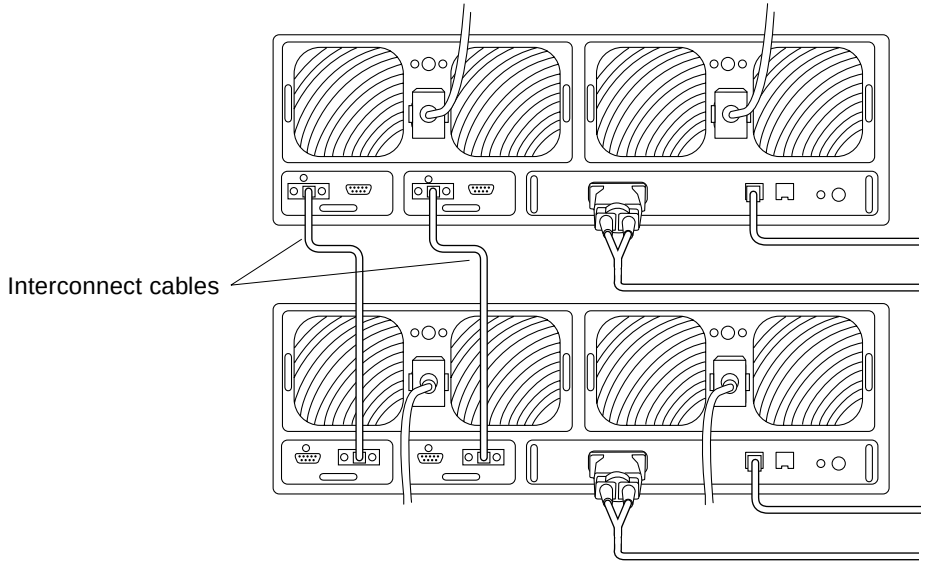

FIGURE 10-4 Interconnect Cable Location

At this point, you can physically move the arrays apart. If you are moving the arrays to different locations, remove the other cables. Replace all cables except the interconnect cable, when the arrays are at their permanent locations.

Note – Do not power on the arrays until you complete the instructions in “Establishing a New IP Address” on page 151.

Establishing a New IP Address

In a partner group, the alternate master unit assumes the IP address of the master unit. When the partner group is disconnected, assign a new IP address to the previous alternate master unit for it to operate as a single controller unit. The JumpStart feature automatically downloads a newly assigned IP address to the array. To enable this feature, you must edit your host file on a RARP server before powering on the array. After you power on, the IP address is automatically assigned. See “Establishing a New IP Address” on page 141 for detailed instructions.

Establishing a Network Connection

After powering on, establish a network connection to each array. This ensures that both arrays function properly and recognize the host.

1. **On the host, use the** `telnet` **command with the array name (or IP address) to connect to the array.**

```
# telnet array_name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

Telnet session (129.150.47.101)
```

Note – The Telnet session verifies that your netwrok connection is good. If the IP address is not assigned correctly, you need to verify the IP adddress in a serial cable connection to make sure that the RARP server is functional.

---

2. **Log in to the array by typing** `root` **and your password at the prompts.**
   - If you are logging in to the previous master unit, use the password for that unit.
   - If you are logging in to the previous alternate master unit, you need to assign a new password. When prompted for a password, press Return.

Note – If you need to create a new password or change some of the parameters, such as the gateway, netmask, and others, refer to Chapter 2 of the *Sun StorEdge T3+ Array Installation and Configuration Manual* for instructions.

---

3. **Use the** `vol list` **and** `vol stat` **commands to verify that the phantom volume has been deleted and that the existing volume remains.**

   For example:

```
:/:<3> vol list
volume          capacity     raid  data      standby
vol2            236.058 GB   5     u2d1-8 none

:/:<4> vol stat
volume          capacity     raid  data      standby
vol2            236.058 GB   5     u2d1-8    u2d9
```

4. **Use the `fru list` and `fru stat` commands to verify that the array is functional and ready for operation.**

   For example:

```
:/:<5> fru list
ID      TYPE                VENDOR       MODEL          REVISION       SERIAL
------  ------------------  -----------  -----------  -------------  --------
u1ctr   controller card     0301         501-5710-02( 020100/020101  112035
u2ctr   controller card     0301         501-5710-02( 020100/020101  112122
u1d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMKJ
u1d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH2A
u1d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1H9WS
u1d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM64
u1d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMC2
u1d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM63
u1d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1HE3A
u1d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNK0
u1d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM5P
u2d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HHH5
u2d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMJC
u2d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1HGKR
u2d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HLBJ
u2d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNH0
u2d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH4Z
u2d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1H92W
u2d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HN9T
u2d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HKQP
u1l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    1413
u1l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    2294
u2l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    001415
u2l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    002054
u1pcu1  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001787
u1pcu2  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001784
u2pcu1  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001544
u2pcu2  power/cooling unit  TECTROL-CAN  300-1454-01( 0000           001545
u1mpn   mid plane           SCI-SJ       370-3990-01- 0000           000953
u2mpn   mid plane           SCI-SJ       370-3990-01- 0000           000958
```

```
:/:<6> fru stat
CTLR    STATUS   STATE       ROLE        PARTNER    TEMP
------  -------  ----------  ----------  -------    ----
u1ctr   ready    enabled     master      u2ctr      31.0
u2ctr   ready    enabled     alt master  u1ctr      30.5

DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      30    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      29    vol1
u1d6    ready    enabled     data disk   ready      ready      30    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      32    vol3
u2d1    ready    enabled     data disk   ready      ready      34    vol2
u2d2    ready    enabled     data disk   ready      ready      38    vol2
u2d3    ready    enabled     data disk   ready      ready      36    vol2
u2d4    ready    enabled     data disk   ready      ready      37    vol2
u2d5    ready    enabled     data disk   ready      ready      34    vol2
u2d6    ready    enabled     data disk   ready      ready      36    vol4
u2d7    ready    enabled     data disk   ready      ready      35    vol4
u2d8    ready    enabled     data disk   ready      ready      40    vol4
u2d9    ready    enabled     data disk   ready      ready      36    vol4

LOOP    STATUS   STATE       MODE        CABLE1     CABLE2     TEMP
------  -------  ----------  -------     ---------  ---------  ----
u2l1    ready    enabled     master      installed  -          29.5
u2l2    ready    enabled     slave       installed  -          31.0
u1l1    ready    enabled     master      -          installed  29.5
u1l2    ready    enabled     slave       -          installed  30.0

POWER   STATUS   STATE       SOURCE  OUTPUT  BATTERY  TEMP    FAN1    FAN2
------  -------  ---------   ------  ------  -------  ------  ------  ------
u1pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u1pcu2  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu2  ready    enabled     line    normal  normal   normal  normal  normal
```

If the array reports a ready status with functional FRUs, you can now restore the data, if necessary, and return the array to operation as a single controller unit.

Alternate Master Unit to a Single Controller Unit

The former alternate master unit might be operating on an outdated file system. If you apply a firmware patch to the arrays while they were connected in the partner group, the files contained on the array's reserved system area are not upgraded on the alternate master, but only on the master unit. When the units are disconnected, the alternate master unit reverts to the file system stored on its reserved system area.

To correct this situation and ensure that the array is ready for operation:

1. **Install the latest firmware patch on the array.**

   This patch is available on the SunSolve web site: `http://sunsolve.sun.com`

   a. **From the SunSolve web site, select Patches under the SunSolve Online column.**

   b. **Select the Storage Products option from the Patches web page.**

      Refer to the `README` file on the web page for specific details on installing the patch for the Sun StorEdge T3+ array firmware.

2. **Create a volume and initialize it.**

3. **Use the `vol list` and `vol stat` commands to verify that the volume(s) is mounted correctly.**

   For example:

```
:/:<7> vol list

volume          capacity     raid  data       standby
vol1            134.890 GB   5     u1d1-5     none

:/:<8> vol stat
vol1           u1d1   u1d2   u1d3   u1d4   u1d5
mounted         0      0      0      0      0
```

4. **Use the `.vol init vol1 fast` command to preserve the old alternate master's data.**

5. **Use the** `fru list` **and** `fru stat` **commands to verify that the array is functional and ready for operation.**

   For example:

```
:/:<9> fru list
ID      TYPE                VENDOR       MODEL          REVISION       SERIAL
------  ------------------  -----------  -----------  -------------  --------
u1ctr   controller card     0301         501-5710-02(  020100/020101  112035
u2ctr   controller card     0301         501-5710-02(  020100/020101  112122
u1d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMKJ
u1d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH2A
u1d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1H9WS
u1d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM64
u1d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMC2
u1d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM63
u1d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1HE3A
u1d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNK0
u1d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HM5P
u2d1    disk drive          SEAGATE      ST336704FSUN A726           3CD1HHH5
u2d2    disk drive          SEAGATE      ST336704FSUN A726           3CD1HMJC
u2d3    disk drive          SEAGATE      ST336704FSUN A726           3CD1HGKR
u2d4    disk drive          SEAGATE      ST336704FSUN A726           3CD1HLBJ
u2d5    disk drive          SEAGATE      ST336704FSUN A726           3CD1HNH0
u2d6    disk drive          SEAGATE      ST336704FSUN A726           3CD1HH4Z
u2d7    disk drive          SEAGATE      ST336704FSUN A726           3CD1H92W
u2d8    disk drive          SEAGATE      ST336704FSUN A726           3CD1HN9T
u2d9    disk drive          SEAGATE      ST336704FSUN A726           3CD1HKQP
u1l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    1413
u1l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    2294
u2l1    loop card           SCI-SJ       375-0085-01-  5.02 Flash    001415
u2l2    loop card           SCI-SJ       375-0085-01-  5.02 Flash    002054
u1pcu1  power/cooling unit  TECTROL-CAN  300-1454-01(  0000          001787
u1pcu2  power/cooling unit  TECTROL-CAN  300-1454-01(  0000          001784
u2pcu1  power/cooling unit  TECTROL-CAN  300-1454-01(  0000          001544
u2pcu2  power/cooling unit  TECTROL-CAN  300-1454-01(  0000          001545
u1mpn   mid plane           SCI-SJ       370-3990-01-  0000          000953
u2mpn   mid plane           SCI-SJ       370-3990-01-  0000          000958
```

```
:/:<10> fru stat
CTLR    STATUS   STATE       ROLE        PARTNER    TEMP
------  -------  ----------  ----------  -------    ----
u1ctr   ready    enabled     master      u2ctr      31.0
u2ctr   ready    enabled     alt master  u1ctr      30.5

DISK    STATUS   STATE       ROLE        PORT1      PORT2      TEMP  VOLUME
------  -------  ----------  ----------  ---------  ---------  ----  ------
u1d1    ready    enabled     data disk   ready      ready      30    vol1
u1d2    ready    enabled     data disk   ready      ready      31    vol1
u1d3    ready    enabled     data disk   ready      ready      30    vol1
u1d4    ready    enabled     data disk   ready      ready      29    vol1
u1d5    ready    enabled     data disk   ready      ready      29    vol1
u1d6    ready    enabled     data disk   ready      ready      29    vol3
u1d7    ready    enabled     data disk   ready      ready      34    vol3
u1d8    ready    enabled     data disk   ready      ready      37    vol3
u1d9    ready    enabled     data disk   ready      ready      31    vol3
u2d1    ready    enabled     data disk   ready      ready      33    vol2
u2d2    ready    enabled     data disk   ready      ready      38    vol2
u2d3    ready    enabled     data disk   ready      ready      36    vol2
u2d4    ready    enabled     data disk   ready      ready      37    vol2
u2d5    ready    enabled     data disk   ready      ready      34    vol2
u2d6    ready    enabled     data disk   ready      ready      36    vol4
u2d7    ready    enabled     data disk   ready      ready      35    vol4
u2d8    ready    enabled     data disk   ready      ready      40    vol4
u2d9    ready    enabled     data disk   ready      ready      36    vol4

LOOP    STATUS   STATE       MODE        CABLE1     CABLE2     TEMP
------  -------  ----------  -------     ---------  ---------  ----
u2l1    ready    enabled     master      installed  -          29.5
u2l2    ready    enabled     slave       installed  -          31.0
u1l1    ready    enabled     master      -          installed  29.5
u1l2    ready    enabled     slave       -          installed  30.5

POWER   STATUS   STATE       SOURCE  OUTPUT  BATTERY  TEMP    FAN1    FAN2
------  -------  ---------   ------  ------  -------  ------  ------  ------
u1pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u1pcu2  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu1  ready    enabled     line    normal  normal   normal  normal  normal
u2pcu2  ready    enabled     line    normal  normal   normal  normal  normal
```

If the array reports a ready status with functional FRUs, you can now restore the data, if necessary, and return the array to operation as a single controller unit.

Changing the Port ID on the Array

To add a partner group to a hub configuration, you must set the port ID values on the arrays to unique values. Sun systems support hard addressing only. However, the `port` command on the Sun StorEdge T3+ array contains the option to set soft addressing. Changing the setting to soft addressing can create problems with host HBAs. In addition, with soft addressing, there is the risk of ending up with new c*x*t*x*d*x* node names after performing a system reboot.

Note – Sun StorEdge T3+ arrays that are factory configured in cabinets with hubs have unique port ID values assigned. This procedure applies only to standalone partner groups that are being added to an existing hub configuration.

To change the port ID on a Sun StorEdge T3+ array:

1. **Connect to the array in a telnet session.**

   See “Establishing a Telnet Session” on page 9 for instructions.

2. **Use the `port` command on the array to change the port ID.**

   You must select a new numerical value for the port identifier. For example, to change a port id on u1p1 from a value of 1 to a value of 20, Type:

```
:/:<1> port set u1p1 targetid 20
```

3. **On the array, type `reset` for the new port ID to take effect.**

APPENDIX **A**

Illustrated Parts Breakdown

This appendix contains part numbers and illustrations of field-replaceable units (FRUs).

The following assemblies are illustrated in this chapter:

Sun StorEdge T3+ Array

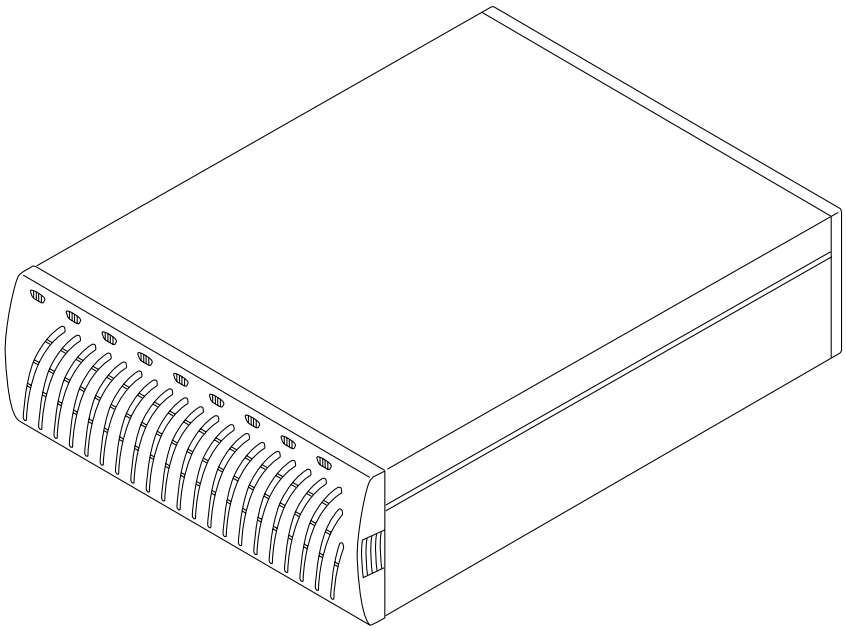

FIGURE A-1 Sun StorEdge T3+ Array (Front View)

Sun StorEdge T3+ Array Assemblies

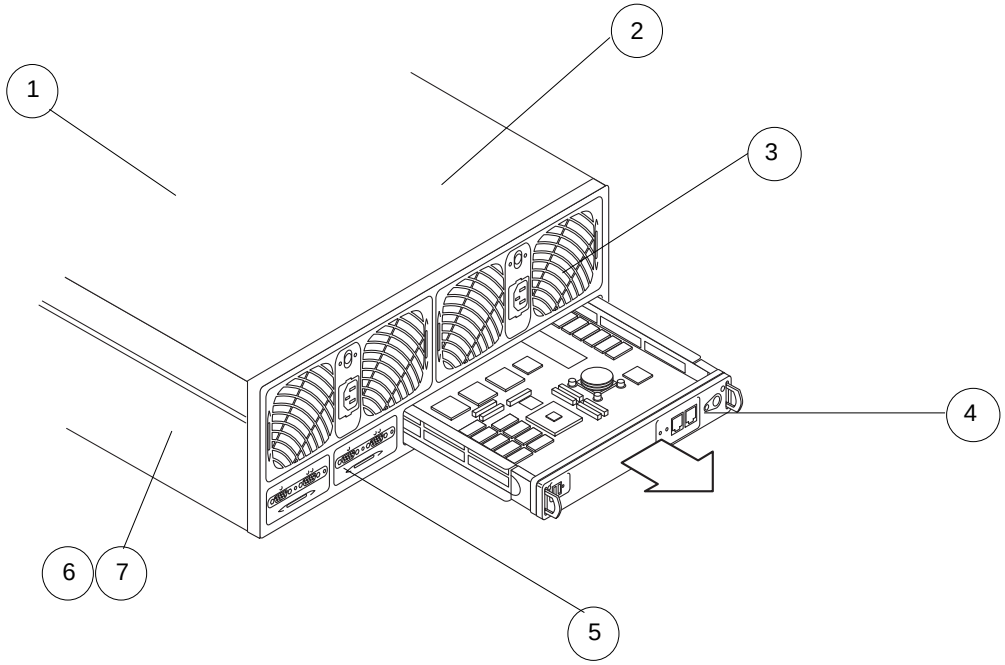

FIGURE A-2 Sun StorEdge T3+ Array (Back View)

TABLE A-1 Sun StorEdge T3+ Array Assemblies

Item	Part Number	Description
1	F540-4306	Door assembly
2	F370-3990	Empty chassis/backplane assembly
3	F300-1454	Power supply and cooling unit
4	F501-5710	T3+ controller card
5	F375-0085	Interconnect card assembly
6	F540-4287	Drive assembly, 18 GB (not shown in this view)
7	F540-4367	Drive assembly, 36 GB (not shown in this view)

Door Assembly

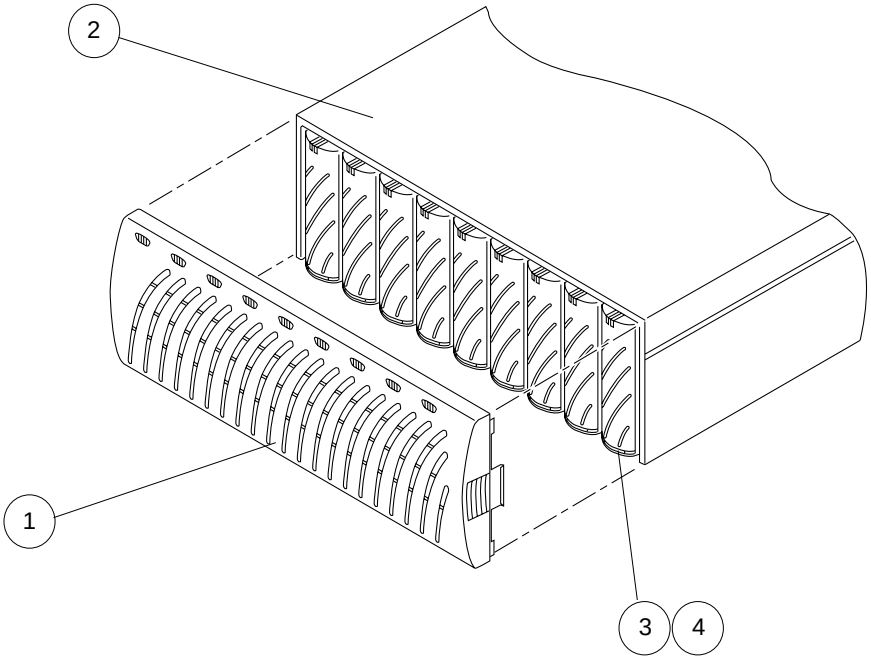

FIGURE A-3 Door Assembly

TABLE A-2 Door Assembly

Item	Part Number	Description
1	F540-4306	Door assembly
2	F370-3990	Empty chassis/backplane assembly
3	F540-4287	Drive assembly, 18 GB
4	F540-4367	Drive assembly, 36 GB

Interconnect Card Assembly

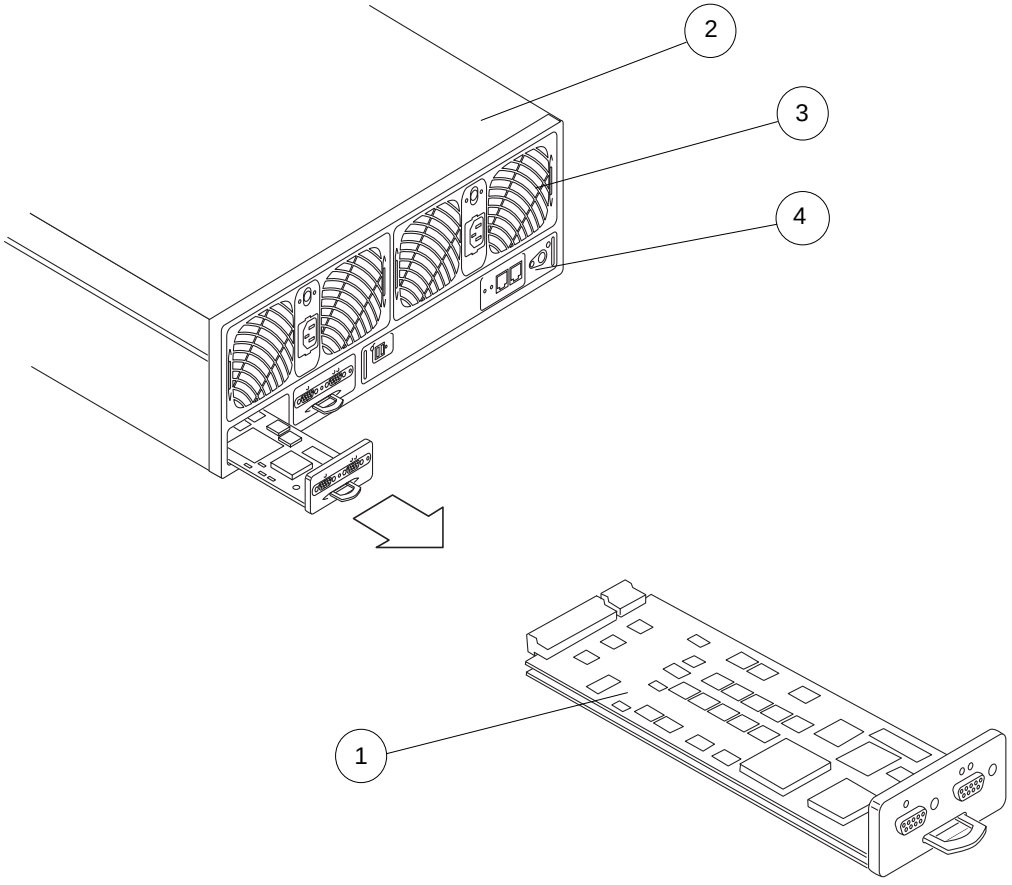

FIGURE A-4 Interconnect Card Assembly

TABLE A-3 Interconnect Card Assembly

Item	Part Number	Description
1	F375-0085	Interconnect card assembly
2	F370-3990	Empty chassis/backplane assembly
3	F300-1454	Power supply and cooling unit
4	F501-5710	T3+ controller card

Power Supply and Cooling Unit

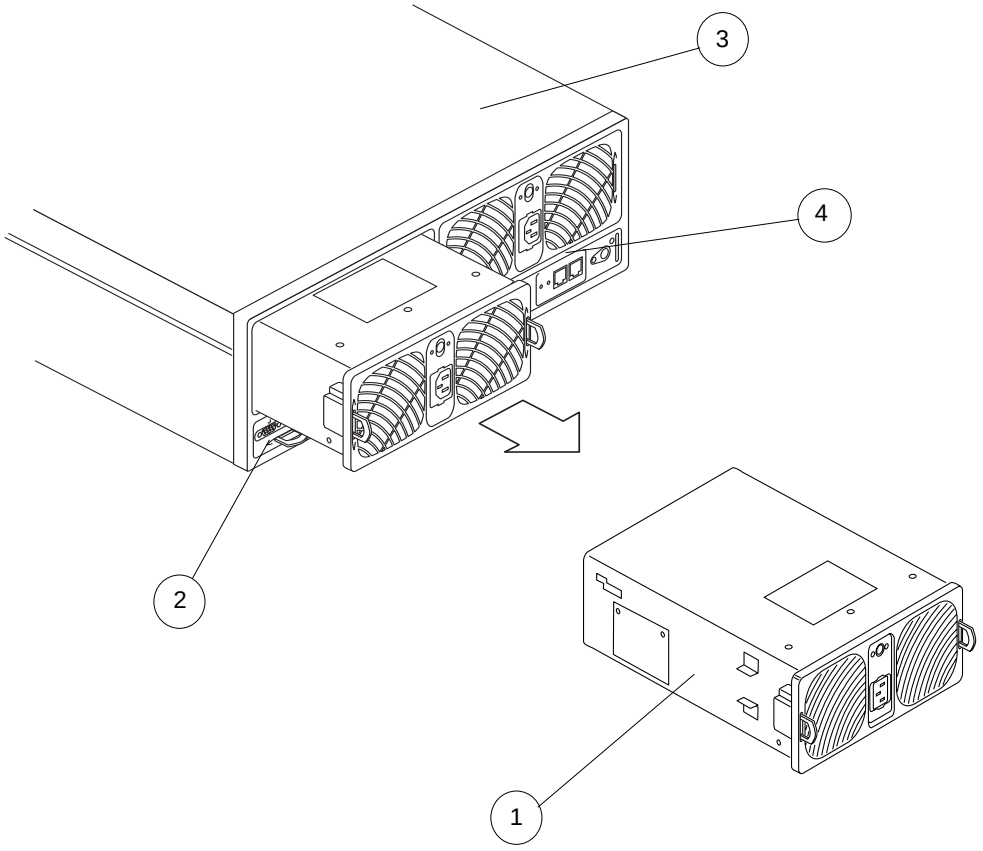

FIGURE A-5 Power Supply

TABLE A-4 Power Supply

Item	Part Number	Description
1	F300-1454	Power supply and cooling unit
2	F375-0085	Interconnect card assembly
3	F370-3990	Empty chassis/backplane assembly
4	F501-5710	T3+ controller card
5	F370-3956	Battery pack, NIMH

Controller Card

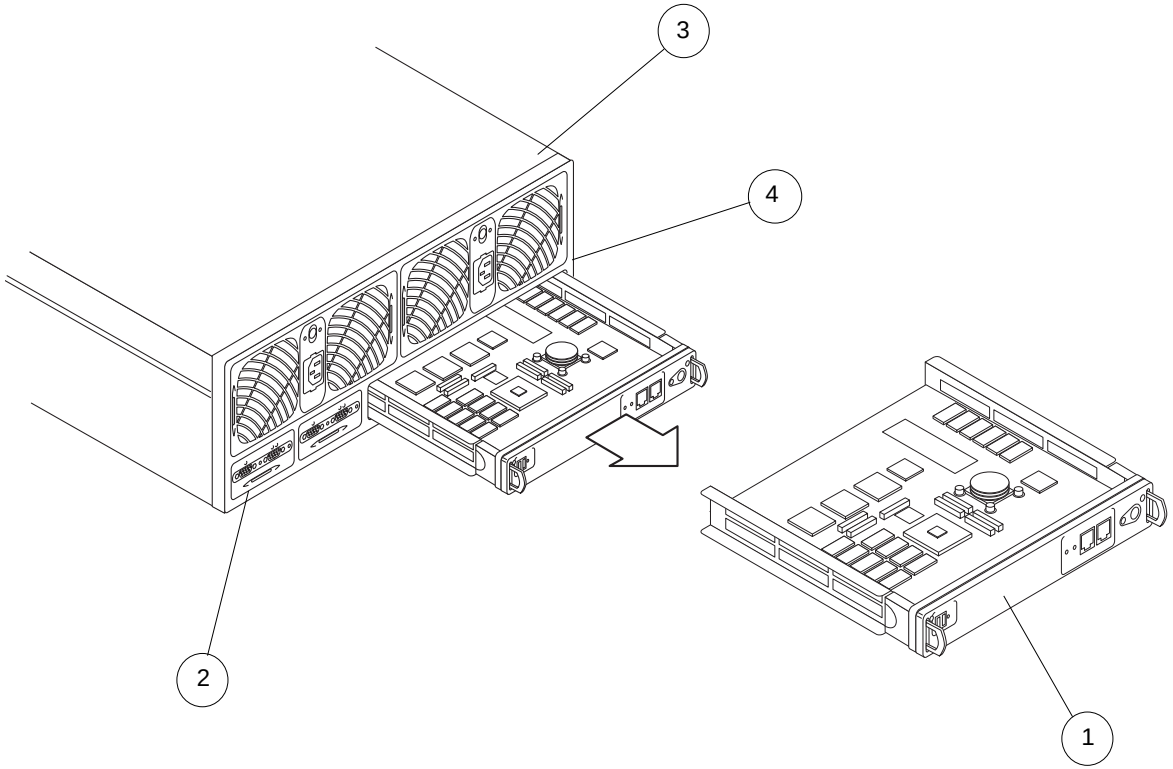

FIGURE A-6 Controller Card

TABLE A-5 Controller Card

Item	Part Number	Description
1	F501-5710	T3+ controller card
2	F375-0085	Interconnect card assembly
3	F370-3990	Empty chassis/backplane assembly
4	F300-1454	Power supply and cooling unit

Drive Assembly

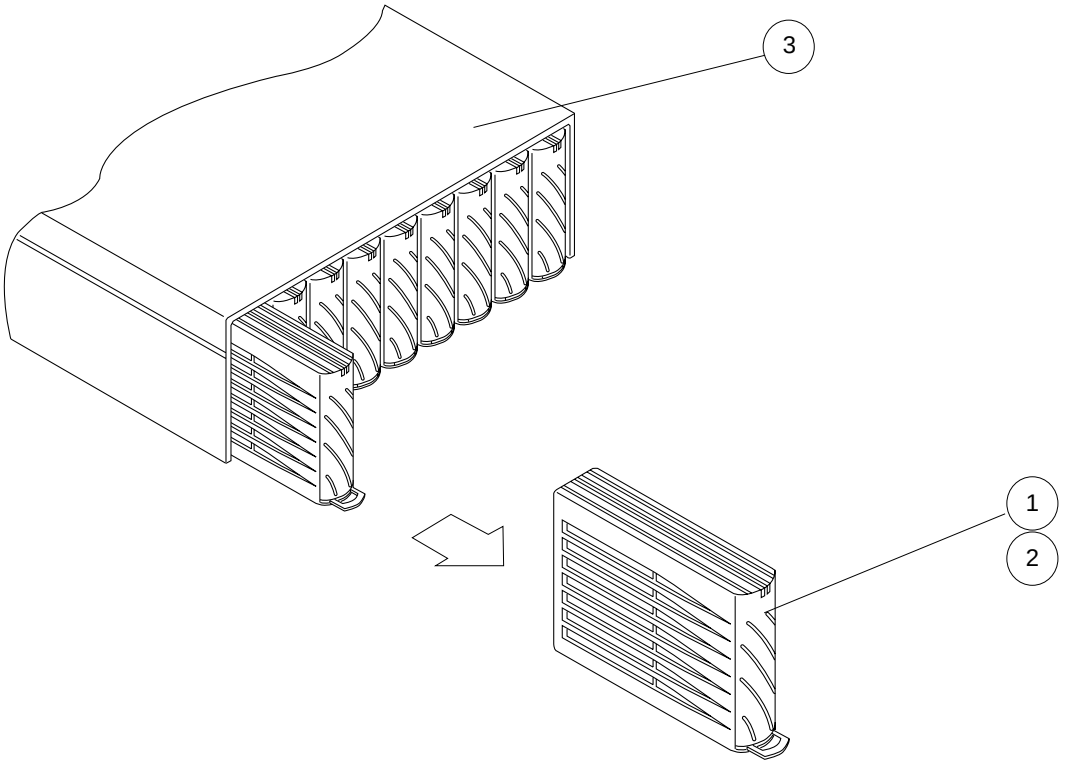

FIGURE A-7 Drive Assembly

TABLE A-6 Drive Assembly

Item	Part Number	Description
1	F540-4287	Drive assembly, 18 GB
2	F540-4367	Drive assembly, 36 GB
3	F370-3990	Empty chassis / backplane assembly

Cable and Interconnect Assemblies

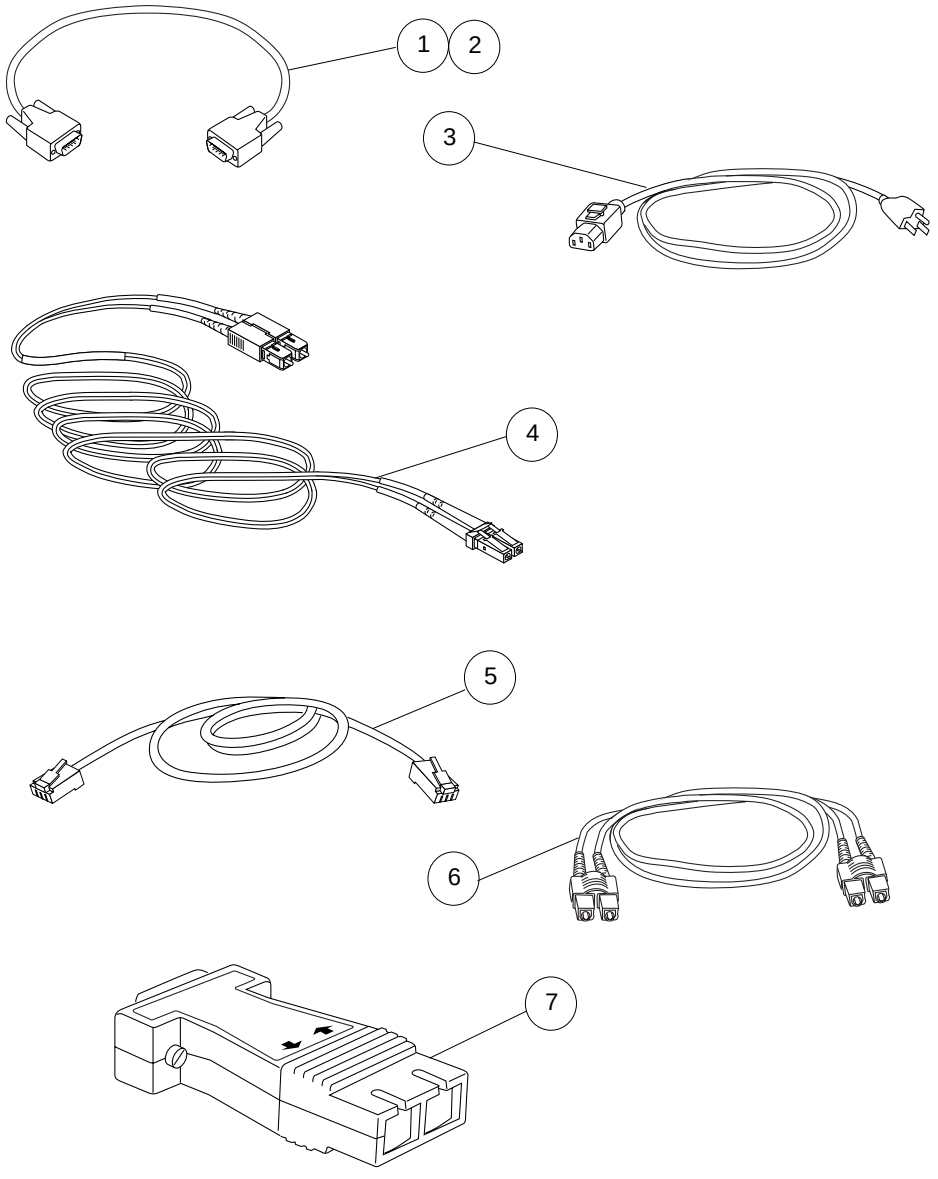

FIGURE A-8 Cables and Interconnects

TABLE A-7 Cable and Interconnect Assemblies

Item	Part Number	Description
1	F530-2842	Interconnect cable, short
2	F530-2843	Interconnect cable, long
3	F180-1918	Locking power cord
4	F537-1034	Fiber-optic cable, Sun StorEdge T3+ array (LC-SFF to SC)
5	[1]	Shielded Ethernet cable (category 5)
6	F537-1020	Fiber-optic cable, Sun StorEdge T3 array
7	F370-3989	MIA adapter

1. Found in F370-4119-02 Diagnostic Cable Kit

APPENDIX **B**

Sun StorEdge T3+ Array System Defaults

This appendix lists the Sun StorEdge T3+ array defaults and is divided into the following sections:

Boot Defaults

Specify boot defaults with the `set` command. When run without any parameters, the `set` command displays the current values. See the *Sun StorEdge T3+ Array Administrator's Manual* for information on using the `set` command.

TABLE B-1 Default Settings—`set` List

Parameter	Default	Variables
`bootmode`	`auto`	[auto | tftp | none]
`bootdelay`	`3`	Number of seconds
`sn`	*Number*	Serial number
`ip`	*n.n.n.n*	Unit IP address
`netmask`	`255.255.255.0`	Unit netmask
`gateway`	*n.n.n.n*	Network gateway IP address
`tftphost`	*n.n.n.n*	IP address of TFTP server

TABLE B-1 Default Settings—`set` List

Parameter	Default	Variables
`tftpfile`	*value*	Boot code file identification number (39 character maximum)
`hostname`	*machinename*	Machine name of the Sun StorEdge T3+ host machine (39 character maximum)
`vendor`	*vendorname*	Name of manufacturer or vendor
`model`	*modelnumber*	Controller model number (set at EP level)
`revision`	`0`*nnn*	Controller EP revision (EP writes this value)
`logto`	*	[1 | * | *filename*] where: 1=Forces logging to serial console *=Directs logging daemon to direct logging as specified in the `/etc/syslog.conf` file.
`loglevel`	`3`	[0|1|2|3|4-] where: 0=No logging at all 1=Error messages only 2=Warning and higher messages 3=Notice and higher messages 4=All message levels including info
`rarp`	`on`	[on | off]
`mac`	*n:n:n:n:n:n*	Controller MAC address. Set by firmware

System Defaults

Specify system defaults with the `sys` command. See the *Sun StorEdge T3+ Array Administrator's Manual* for more information on using the `sys` command.

TABLE B-2 System Default Settings

Sys Parameter	Default	Variables
`blocksize`	`64k`	[16k | 32k | 64k]
`cache`	`auto`	[auto|writebehind|writethrough|off]
`mirror`	`auto`	[auto | off]
`mp_support`	`none`	[rw | none] Multi-pathing support
`rd_ahead`	on	Set to off to always perform datablock read ahead.

TABLE B-2 System Default Settings

Sys Parameter	Default	Variables
`recon_rate`	medium	[high, medium, low] Reconstruction rate.
`memsize`	32	Set by controller, read-only. In MBytes
`cache memsize`	256	Set by controller, read-only. In MBytes

Volume Defaults

Specify system defaults with the `vol` command. See the *Sun StorEdge T3+ Array Administrator's Manual* for more information.

TABLE B-3 Volume Defaults

Parameter	Default	Variables
`init [rate` *n*`]`	16	[1-16] 1 is lowest, 16 is highest...
`verify [rate` *n*`]`	1	[1-8] Rate parameter refers to host interleave factor (contention with host IOs). Default is 1. There is currently no feature that spawns a `vol verify` process.

The default for the SCSI "vendor ID field" is `Sun`. Display or change this value with the `port` command.

The default Sun StorEdge T3+ array volume configuration as shipped from the factory is 9 drive RAID 5 with no standby disk.

The volume is configured as follows:

- vol add v0 data u1d1-9 RAID 5
- vol init v0 sysarea
- vol init v0 data rate 16

Default Directories and Files

TABLE B-4 lists the default file system shipped with the array.

TABLE B-4 Default Directories and Files

Filename	Description
`/nb113.bin`	Controller firmware (RR Sum: 23020 5000)
`/lplc_05.01`	Interconnect card FW (RR Sum: 63295 21)
`/BITMAP.SYS`	Contains a map of used and free blocks
`/ep2_10.bin`	Controller EPROM flash (RR Sum: 3221 1023)
`/FLIST.SYS`	Contains the file descriptors
`/cmdlog`	Log of all commands executed on the system
`/adm`	Legacy directory formerly used for syslog files
`/webgui`	Contains old browser based admin files
`/etc/hosts`	Default hosts with comments on format of file
`/etc/schd.conf`	Battery refresh file. Contents: BAT_CYC 14
`/etc/syslog.conf`	System logging configuration file
`/syslog`	Default system logging file
`/web/snmp/T3.mib`	SNMP required file

Note – At time of manufacture, Sun StorEdge T3+ array system disks do not contain controller firmware, interconnect card binaries, EP binaries, or drive firmware images. You can download all of these from the SunSolve web site.

APPENDIX **C**

Sun StorEdge T3+ Array Messages

This appendix contains a description of array error messages. See the *Sun StorEdge T3+ Array Administrator's Manual* for explanations of the more important error messages. This chapter contains the following sections:

Message Syntax

Error message syntax consists of the following two components:

- "Message Types" on page 174
- "FRU Identifiers" on page 175

The following sub-sections describe these components and list possible error and warning messages. See the *Sun StorEdge T3+ Array Administrator's Manual* for explanations of the more important error messages.

Message Types

A syslog daemon exists in the hardware RAID controller that records system messages and provides for remote monitoring. There are four levels of messages, listed in TABLE C-1 in order of severity. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* to use the `set` command to set the `loglevel` to receive notification of the various types of messages.

TABLE C-1 Message Types

Message Type	Definition
Error	Indicates a critical system event requiring immediate user intervention or attention. For example, an over temperature condition or a detected FRU being removed.
Warning	Indicates a possible event requiring eventual user intervention. For example, a FRU being disabled and recovery procedure executed.
Notice	Indicates a system event that might be a side effect of other events or may be a normal condition. For example, the power switch is turned off.
Information	Indicates a system event that has no consequence on the running health of the system. For example, a good state of a FRU.

FRU Identifiers

The syntax of the error message uses a FRU identifier to refer to a particular FRU in a disk tray. This identifier contains a unit constant (u), the unit number (*n*), the FRU constant (ctr for controller card, pcu for power and cooling unit, 1 for interconnect card, d for disk drive), and the FRU number (*n*).

TABLE C-2 FRU Identifers

FRU	Identifier	Unit Number
Controller card	u*n*ctr	*n* = unit number (1, 2, ...)
Power and cooling unit	u*n*pcu*n*	*n* = unit number (1, 2, ...) *n* = pcu number (1, 2)
Interconnect card	u*n*l*n*	*n* = unit number (1, 2, ...) *n* = interconnect number (1, 2)
Disk drive	u*n*d*n*	*n* = unit number (1, 2, ...) *n* = disk drive number (1 to 9)

Miscellaneous Abbreviations

LPC: Loop card

BATD: Battery monitor

IPI-3: Intelligent Peripheral Interface.Similar legacy protocol to SCSI, the Sun StorEdge T3+ array uses IPI-3 for configuration data

TDL: Transaction disk log

CCB: Command Control Block

SCB: Stripe Control Block

IOCB: ISP2100 IO Control Block. Basically a request put into the queue for the ISP to process

IOSB: ISP2100 Status Block

SVD: SCSI Virtual Disk Driver. This driver is the backend disk driver in the T3

SVH: SCSI virtual host driver. The front-end Sun StorEdge T3+ array driver which takes host requests for ISP2100 in target mode.

XPT: SCSI Transport Layer module in Sun StorEdge T3+ array driver stack

SID: Stripe ID

STYPE: Stripe type

ISR: Interrupt service routine

Interpreting Sun StorEdge T3+ Array /syslog Messages

When attempting to sort out Sun StorEdge T3+ array related loop problems, or interpret the Sun StorEdge T3+ array syslog, it is important to have the data host messages file available. If you are troubleshooting a live array, you should always enable remote syslogging and monitor the host messages and array messages at the same time. A laptop and ethernet hub come in handy here. It is also important to verify that the time and date are the same on both the arrays and the data hosts.

The following commands are also useful for finding all the targets and initiators on the loop in question:

```
# luxadm -e port
```

```
# luxadm -e dump_map device
```

where *device* is from the output of the previous command.

To find the targetids and WWNs of the array ports, use:

```
:/:<n> port list
```

```
:/:<n> port listmap
```

There is a table of internal alpa/target mapping in the appendix of this document.

Armed with this information you should be able to sort out who is who, and what is what.

When debugging, it is also useful to reset the syslog on the array and the remote syslog host to clear out any noise from earlier testing, problems, or the initial install:

```
:/:<n> set logto 1
```

```
:/:<n> mv syslog syslog.bak
```

```
:/:<n> logger "Starting New Syslog xx/xx/xxxx" > syslog
```

```
:/:<n> set logto *
```

The Basic Message Format

```
May 18 16:36:08 FCC0[1]: N: u1ctr (ITL 7D 1 0 TT 20 TID 9684 OP 4D) Invalid
command opcode
```

where:

- `May 18 16:36:08` = date and time
- `FCC0` = the task that generates the message
- `[1] u1ctr` = the controller reporting the error
- `N` = message level
- `u1ctr` = FRU identifier
- `(ITL 7D 1 0 TT 20 TID 9684 OP 4D) Invalid command opcode` = message text

The first thing to look at is the task. There is a list of tasks in the FAQ on the HES website and at the end of this document. The most important information for a quick look at data path problems or LUN access problems is the task. If you see “FCC0” you know immediately this is a host port issue and you probably have a front end

loop problem. “FCC2” is the cache mirroring task, These represent chatter between the controllers to monitor the status of each others cache mirror. The FCC2 messages can be misleading since the cache mirror is actually seen as a LUN, which means you get messages just like the one above. But the LUN being queried is a “virtual” LUN. These are typically seen right after a boot or when the cache is being flushed (see explanation below).

There are four levels of messages, listed in TABLE C-1 in order of severity: [E]rror, [W]arn, [N]otice, and [I]nformation. Be careful to observe all [E]rrors. These are critical events like FRU failures. [W]arnings are important as well and could indicate a future problem. [N]otices are frequent and voluminous. Many are just “chatter” between the controllers on loop 2 (when ever a cache flush occurs), common system events like a battery refresh or reporting the temperature of the loop cards, and host related events like reboots etc. However, an error or warning will often trigger a cascade of notice messages indicating LUN takeovers, cache flushing, etc. Like LIPs on an FCAL, a few are ok and expected, but you should pay attention to storms of them, or frequent repeats of the same message. Note that they often contain useful debugging information that can help determine the root cause of a failure. When you see these patterns, look back in the syslog for the [E]rror or [W]arning that precipitated them.

Interpreting ITL Messages in an FCAL Environment

(I)nitiator, (T)arget, (L)UN messages record SCSI commands being received by the various port monitoring tasks. They are common and frequently not a cause for concern.

Basic Example

```
(ITL 7D 5 1 TT 20 TID 9CA8 OP 0) Target in Unit Attention
```

where:

- ITL = (I)nitiator, (T)arget, (L)UN
- TT = Tag Type = 20 and a tag type of 20 is a Simple Queue Tag
- TID = Tag ID = 9CA8 which is the unique (for this I/O) Tag Id number.
- OP Code = SCSI OP code = 0 which is "Test Unit Ready"

Note – The initiator can be verified using “luxadm -e dump_map <device>.” Other common initiator HBA ID’s are: 7C = dec 124 and 7B = dec 123

So the things that are probably going to be most useful in field based diagnosis are; the Target (the drive), the LUN, and the OP code, which will generate either a response, or the actual OP code text itself.

ITL Message Examples

Host Port Message

```
FCC0[1]: N: u1ctr (ITL 7D 5 1 TT 20 TID 9CA8 OP 0) Target in Unit Attention
```

where:

- FCC0 = task on external loop 3 - host loop
- [1] = enclosure_id 1 (u1)
- 7D = initiator 7D (alpa x01) HBA on host
- 5 = target 5 (alpa xEF)
- 1 = LUN 1

This is a very common message. Seen during the Sun StorEdge T3+ array boot sequence or as the result of host activity such as a reboot and luxadm inquiry. The initiator is sending a SCSI command to LUN 1 on controller 1.

Back-End Loop Message

```
FCC2[1]: N: u1ctr (ITL 1 0 1 TT 20 TID AAE8 OP 0) Target in Unit Attention
```

where:

- FCC2 = task on loop 2
- [1] = enclosure_id 1 (u1)
- 1 = initiator 1 (alpa xE8), ISP chip on ctrl u2
- 0 = target 0(alpa xEF) ISP chip on ctrl u1
- 1 = LUN 1

In this case ctrl2 (u2) sent a Test Unit Ready cmd (OP 0) through loop 2 and u1 responds with Unit Attention. u2 is checking on the status of the cache mirroring LUN. We know this because it is task FCC2 and an initiator on a host side loop would have one of the standard initiator AL_PAs like 7C, or 7D like the example above.

Interpreting ITL Messages in a Fabric/SAN Environment

Everything is the same as the FCAL environment except for 2 things you need to be aware of.

1. Beginning with the 1.17 and 2.0 bootcode releases the target is now reported using the hex version of the 7-bit loop ID (The SEL_ID column in the AL_PA chart).
2. You will see initiators with very low numbers like EF and E8. These are fabric ports on a switch and/or 3rd Party HBAs like JNI etc. (check your task - now both FCC0 and FCC2 events can have low initiator numbers)

Port Event Messages

These are typically port login/logout events. Common on the backend when a LUN fails over and on the host side when a host reboots or the loop goes down for some reason.

This is a common host port sequence:

```
ISR1[1]: N: u1ctr ISP2100[2] Received LIP(f7,f7) async event
FCC0[1]: N: u1ctr Port event received on port 0, abort 0 (id 123)
FCC0[1]: N: u1ctr Port event received on port 0, abort 0 (id 124)
ISR1[1]: N: u1ctr ISP2100[2] Received LIP(f7,f7) async event
FCC0[1]: N: u1ctr Port event received on port 0, abort 0 (id 124)
FCC0[1]: N: u1ctr Port event received on port 0, abort 0 (id 123)
```

where `id` is the Initiator.

This Sun StorEdge T3+ array is connected to a loop with 2 initiators - A LIP is received on the host port on u1 and the HBAs/Initiators connected to that port logout and log back in. You would see something similar if a switch port were reset but the “id” would be low on the chart, an E8 for example.

Identifying Sun StorEdge T3+ Array Ports and Loops

Ports

- On a single Sun StorEdge T3+ array there are 3 ports; 2 backend, 1 host port.
- On a partner pair there are 6; 4 backend, 2 host ports.

Loops

- FW 1.17 and older: 2 internal loops + 1 external host loop.
- FW 1.18/2.0+: 3 internal loops + 1 external loop.

So, for a T3PP:

- loop1 (path_id=0) connects ports 1 and 4
- loop2 (path_id=1) connects ports 2 and 5
- host ports are 0(u1),3(u2)

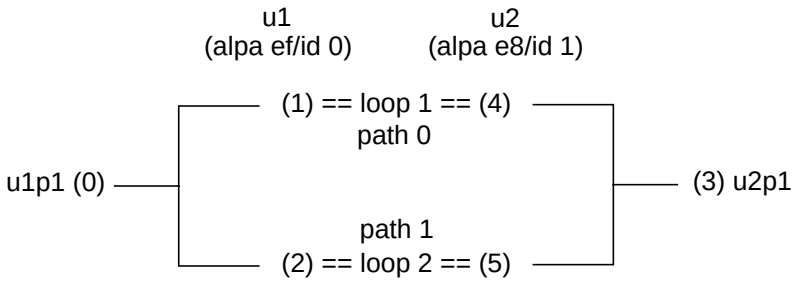

FIGURE C-1 Loop/Port Diagram

Note – You will see references to “ports” in different contexts. Although the above information is accurate, each disk is also a port on the backend loops but is referenced differently. They are referenced in the context o f SCSI errors and are identified with a hex number which corresponds to the “SEL_ID” column in TABLE C-12.

So, FCC2[2]: N: u2ctr port event received on port 5, abort 0 (id 0). In this case ISP port 2 (initiator id 0, loop 2) on ctrl1 did a login or logout generating on ISP port 5 on ctrl 2 a ‘Port event.’

Note – Regarding these messages: You should only see chatter between ports 1 and 4 when loop 2 has failed and loop 1 is healthy.

Backend Loop chatter - Loop 2 - cache mirroring

Jun 02 05:41:34 FCC2[1]: N: u1ctr Port event received on port 2, abort 0 (id 1)

Jun 02 05:41:36 FCC2[1]: N: u1ctr (ITL 1 0 0 TT 20 TID A308 OP 2A) Target in Unit Attention

Jun 02 05:41:36 FCC2[1]: N: u1ctr <<Abort Task Set>> on port 2, abort 0

Jun 02 05:42:05 FCC2[1]: N: u1ctr Port event received on port 2, abort 0 (id 1)

Jun 02 05:42:07 FCC2[1]: N: u1ctr Port event received on port 2, abort 0 (id 1)

Jun 02 05:42:07 FCC2[1]: N: u1ctr (ITL 1 0 0 TT 20 TID A50C OP 2A) Target in Unit Attention

Jun 02 05:42:07 FCC2[1]: N: u1ctr <<Abort Task Set>> on port 2, abort 0

Jun 02 05:42:07 FCC2[2]: N: u2ctr Port event received on port 5, abort 0 (id 0)

Jun 02 05:42:07 FCC2[2]: N: u2ctr Port event received on port 5, abort 0 (id 0)

Jun 02 05:42:07 FCC2[2]: N: u2ctr Port event received on port 5, abort 0 (id 0)

Jun 02 05:42:08 FCC2[2]: N: u2ctr (ITL 0 1 1 TT 20 TID A6EC OP 2A) Target in Unit Attention

Jun 02 05:42:08 FCC2[2]: N: u2ctr <<Abort Task Set>> on port 5, abort 0

SVD/SVC Error Messages

```
ISR1[1]:W:SVC_PATH_FAILOVER:path_id=0, lid=15
```

where:

- path_id=0, backend loop 1
- lid=logical unit identification (15 = u2d7)
- SVD talks in terms of lid's (LUN id). See TABLE C-3.

TABLE C-3 LIDs corresponding to LUN IDs example

LID	target_id	LUN
1-8 u1d1-8	08 - 0f	0
9-16 u2d1-8	10 - 17	0
17 u1d9	97	0
18 u2d9	98	0
19 v0	1	0
19 v1	0	1
20 v2	1	2
20 v3	0	3

where v0, v1, v2, v3 are volumes created in this order:

v0 - on u1

v1 - on u2

v2 - on u1

v3 - on u2

The lid=19 and 20 are the lun_id assigned to the target cache mirroring LUN, which is a "virtual LUN" that receives SCSI commands just like a real LUN. Hence the variety of the aborted tasks and as seen above. Therefore, it is a "shared" resource and each controller holds a pointer to the virtual LUN representing the stripe set for the other controller's volume. And each volume has it's own stripe set in cache.

In summary, the targets of the volumes are:

LID	Name	target_id	LUN
19	v0	1	0
19	v1	0	1
20	v2	1	2
20	v3	0	3

However the controller targets are:

- u1 = 0
- u2 = 1

The targets appear to be reversed because the mirror for u2v1 actually resides on u1. So lid 19 for v1 has to have a target of 0 since to reach it, an access u1 is needed. Therefore, to eliminate confusion, always try to stick with the following convention for volume creation:

- LUN 0 -> u1
- LUN 1 -> u2
- LUN 2 -> u1
- LUN 3 -> u2

Fatal Timeouts

```
ISR1[2]: N: u2ctr ISP2100[1] Fatal timeout on target 0.1
ISR1[2]: N: u2ctr ISP2100[1] QLCF_ABORT_ALL_CMDS: Command Timeout Pre-Gauntlet
Initiated
```

where:

`target 0.1` refers to portid/target=0 and lun=1.

The translation/address_resolution to the FC drives is done at XPT/SIM level and is using target.lun format, i.e., in this case target 0.1, 0 is target_id, 1 is LUN.

This example shows a fatal timeout on the cache mirroring LUN for LUN 1 on u2. This is evident by using the information from the previous discussion of lids, the ALPA chart, and port listmap.

We now know two things:

- portid / target 0 is the cache mirror LUN 1 on u2
- LUN 1 is on u2

The Gauntlet:

When a controller issues a command to its partner, it starts a watchdog timer for the command. And if the command is not complete within the required time frame, the controller will timeout the command as a fatal error.

This indicates that u2 tried to write to its cache mirror and couldn't. Therefore, it timed out the command. This example is from a case where u1 had failed and was eventually replaced.

Disk-Related Error Messages

Disk drives have CRC and ECC protection on all sectors so they can detect whether or not data is read correctly and in some cases use the ECC to correct the data.

Many disk errors consist of more than a single `syslog` entry. Frequently the event occurring on, with, or to the disk will generate other system events such as a PATH failover, or the disabling of the disk after too many errors. The key is to look for is clusters of messages. After a certain threshold, the active controller disables the failing drive.

Basic Format of Messages

```
ISR1[1]: W: SCSI Disk Error occurred (path=0x0, port=0xc, lun=0x0)
```

where:

- `SCSI Disk Error occurred =` Notification of an event
- `0x0 =` Path event detected on
- `0xc =` Disk identified as having the error. SEL_ID column of the internal AL_PA chart or :/:<1> .sim -f num 0 id2alpa 0xc pass -> [loopid, alpa] -> [0xc, 0xd3]
- `0x0 =` LUN which disk is a part of

```
ISR1[2]: N: u2d1 SCSI Disk Error Occurred (path = 0x1)
```

where:

- `u2d1 =` Path error detected on
- `SCSI Disk Error Occurred =` Notification of an event
- `0x1 =` Disk where error is occurring

RAID Stripe:

```
ISR1[2]: N: u2d8 sid 234096 stype 2023 disk error 3
```

where:

- `sid 234096 =` RAID stripe in cache
- `stype 2023 =` RAID stripe type (see table)
- `error 3 =` Specific error type (see table)

SCSI Disk Errors

These events are recorded by a sequence of 4 messages describing the disk having the error, the path the error is detected on, the actual error, a translation, and the "Valid Information" field. The 1st and 3rd lines are the most important since they tell us which disk had the error and what that error was.

1: 09:58:43 ISR1[1]: N: u1d3 SCSI Disk Error Occurred (path = 0x1)

2: 09:58:43 ISR1[1]: N: Sense Key = 0x1, Asc = 0x17, Ascq = 0x1

3: 09:58:43 ISR1[1]: N: Sense Data Description = Recovered Data With Retries

4: 09:58:43 ISR1[1]: N: Valid Information = 0x26af795

Line 1 - Tells us an error occurred and on what disk

Line 2 - A detailed description of the error reported. See "Sense Key Explanations" on page 211 and the following web site to decipher these.

`http://www.t10.org/lists/1spc-lst.htm`

Line 3 - An explanation of the sense key (see list below).

Line 4 - Not useful

Examples

Recoverable

09:58:43 ISR1[1]: N: u1d3 SCSI Disk Error Occurred (path = 0x1)

09:58:43 ISR1[1]: N: Sense Key = 0x1, Asc = 0x17, Ascq = 0x1

09:58:43 ISR1[1]: N: Sense Data Description = Recovered Data With Retries

09:58:43 ISR1[1]: N: Valid Information = 0x26af795

09:58:58 ISR1[1]: N: u1d3 SCSI Disk Error Occurred (path = 0x1)

09:58:58 ISR1[1]: N: Sense Key = 0x1, Asc = 0x18, Ascq = 0x2

09:58:58 ISR1[1]: N: Sense Data Description = Recovered Data - Data Auto-Reallocated

09:58:58 ISR1[1]: N: Valid Information = 0x26af795

The errors above indicate that the drive had a problem and was able to resolve it by the drive itself re-reading the information and marking a sector bad and auto reallocating the data to an alternate sector.

Parity Errors

12:39:06 ISR1[2]: W: u2d6 SCSI Disk Error Occurred (path = 0x0)

12:39:06 ISR1[2]: W: Sense Key = 0xb, Asc = 0x47, Ascq = 0x0

12:39:06 ISR1[2]: W: Sense Data Description = SCSI Parity Error

12:39:06 ISR1[2]: W: Valid Information = 0x3379602

Common Host Port (FCC0) Messages

```
13:42:41 FCC0[1]: N: u1ctr IDE received on port 0, abort 0
```

where:

`IDE`: Initiator Detected Error

The above message is printed when the host port task. FCC0 receives an abort command from the initiator. The initiator sends the abort when it detects an error on the target. In this case the Sun StorEdge T3+ array LUN being accessed on port 0 (see “Identifying Sun StorEdge T3+ Array Ports and Loops” on page 181).

Check the host syslog. You should see SCSI resets and retries that occurred at the same time.

Assertion and Exception Reset Messages

These occur for one of 2 reasons, a hardware fault generating an “exception,” or a controller encounters an area of code designed to generate an “assertion” in certain scenarios. They are somewhat analogous to a *kernel panic* in solaris. An event has occurred, or situation arises that could result in writing or reading bad data; the controller then *panics*. In a properly configured and healthy enterprise configuration this is no problem. The other controller takes over the LUNs and disables the one experiencing the event. In a workgroup configuration, it resets and you lose access to your LUNs (this presents a good argument for host-based mirroring).

Note – *Important information regarding where the reset is recorded.* Each controller has a space in NVRAM where the last reset is stored. When another exception or assertion reset occurs, it will replace the information in NVRAM with the new information. This information also follows the controller when it is replaced. It is therefore important to capture this information when returning a controller for CPAS. It is also a good idea to clear the reset log on the new controller. See “Reset Log Message Types” on page 191.

Note – *Regarding an enterprise configuration.* The reset you see in the syslog, from extractor, or a reboot, is from the active master controller. To dump the log on the alternate you will need to use the serial port and run the commands from there.

Assertion and 2004 Exceptions are software related. 2003 Exceptions are hardware related. However, you can get an Assertion when a FRU fails, causing a retry threshold on RAID reads to be exceeded.

The useful information is the first line of an assertion and exception. It indicates what type of Reset occurred. This provides an idea of how to proceed. Without access to the source code these messages are almost no value other than indicators that something happened. They must always be interpreted in the context of other events.

- SysFail cache parity reset, replace the controller.

- Assertion - look at the source code, go to the line in the file referenced and determine, based on syslog events whether it is relevant.
- Exception - hard to say since there no source to reference. In these cases you have to wing it. Are there any failed FRUs?

Examples

Cache Parity Error - Replace Ctlr

15:18:35 t3a pshc[1]: W: u1ctr SysFail Reset (7001) was initiated at 20010626

163740 Cache memory parity error detected

Assertion Reset

14:47:16 sh05[1]: W: u1ctr Assertion Reset (3000) was initiated at 20020308 213140

../../common/msc/sxf_task.c line 763, Assert (err == 0) => 0 BOOT

14:47:16 sh05[1]: N: CPU state...

14:47:16 sh05[1]: N: R0 = 000c9ea4 019cf510 002936bc 00000001 00000002 019cf3d0 016408e0 00000001

14:47:16 sh05[1]: N: R8 = 00000001 000000c8 000000c8 004d0000 004cd1a0 00294dec 00000000 00000000

14:47:16 sh05[1]: N: R16= 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000

14:47:16 sh05[1]: N: R24= 0027ad48 0027a900 00000000 00409ef4 00000000 00000000 008fb408 008fb048

14:47:16 sh05[1]: N: CR=40000000 XER=00000000 LR=000c9eec CTR=00000000 DSISR=00000000

14:47:16 sh05[1]: N: DAR=00000000 MSR=0000b930 IP=SRR0=001888ec SRR1= 0000b930

Exception Reset

19:31:53 pshc[1]: W: u1ctr Exception Reset (2004) was initiated at 20010904 192859 Instruction Access exception

19:31:53 pshc[1]: N: CPU state...

19:31:53 pshc[1]: N: R0 = 0008f640 018b57a8 002936bc 00000019 01870000 0164dfe8 018b5e4c 001b6d2c

19:31:53 pshc[1]: N: R8 = 0000b930 0164dfe8 01640d04 004d0000 004cd1a0 00294dec 00000000 00000000

19:31:53 pshc[1]: N: R16= 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000

19:31:53 pshc[1]: N: R24= 00000000 00000000 00000000 00000000 00000000 0008f640 00010400 000004c8

19:31:53 pshc[1]: N: CR=44000000 XER=00000000 LR=0008f650 CTR=00000000 DSISR=00000000

19:31:53 pshc[1]: N: DAR=00000000 MSR=00001030 IP=SRR0=deaddeac SRR1= 4000b930

"DATA LENGTH INCORRECT" from bug id 4355112

04:58:17 FCC2[2]: N: u2ctr Port event received on port 5, abort 0 (id 0)

04:58:17 FCC2[2]: N: u2ctr (ITL 0 1 0 TT 20 TID AB84 OP 2A) Target in Unit Attention

04:58:17 FCC2[2]: N: u2ctr (ITL 0 1 0 TT 20 TID AB90 OP 2A) Aborted by Host

04:58:17 SX01[2]: N: u2ctr (ITL 0 1 0 TT 20 TID AB90 OP 2A) Data length incorrect

04:58:17 FCC2[2]: N: u2ctr <<Abort Task Set>> on port 5, abort 1

All Sun StorEdge T3+ arrays LUNs have a Power-On Unit Attention pending on each port for each initiator. Therefore, the back-end cache mirroring LUN will receive this error condition for the first I/O. Since the SVD disk driver causes a force flush by issuing an Abort Task Set upon receiving a Unit Attention condition, all outstanding cache mirroring LUN commands, at the time of the Unit Attention condition is received, will be aborted. In addition, potential Notice syslog messages may be generated due to the command prematurely getting aborted (for example, if the data length is incorrect).

Note – Once this initial Unit Attention condition is cleared, any subsequent Unit Attention conditions (causing Abort Task Set to be generated) during normal operation may be due to faulty hardware (and is *not* deemed to be “normal and expected behavior”).

Reset Log Message Types

If the `set` command is used with the `loglevel` parameter to set the notification level to “2” (warning and error messages), or higher (3 or 4), you can trace the reason for the reset be examining the contents of the `/syslog` file. This is possible because the reset log information is downloaded into the `/syslog` file every time the system resets. If desired, the reset log information can also be downloaded whenever the `logger -dmprstlog` command is issued.

TABLE C-4 Reset Log Message Types

Index	Type	Type Value	Description
0	Hardware	0x1000	User reset
1	Exception	0x2000	Exception
2	Assertion	0x3000	Software assertion
3	RAIDFail	0x4000	RAID fatal error
4	Takeover	0x5000	Takeover
5	pSOSFail	0x6000	pSOS fatal error
6	SysFail	0x7000	System error

Type the following to capture the log:

```
t3:/:<n> logger -dmprstlog
```

Type the following to clear the log:

```
t3:/:<n> logger -clrrstlog
```

Reset Log Messages

TABLE C-5 Reset Log Messages

Type	Mask	Description
RESET_FAIL	1000	Hardware Reset
EXCEPT_FAIL	2000	--
	2003	Data access exception
	2004	Instruction access exception
	2005	Alignment exception (operand not word aligned)
	2008	Floating Point exception
ASSERT_FAIL	3000	Software detected fault
RAID_FAIL	4000	--
SNXF_IN	4001	Short non-transfer in
SNXF_EXE	4002	Short non-transfer execution (mode sense/select)
SNXF_OUT	4003	Short non-transfer out
LNXF_IN	4004	Long non-transfer in
LNXF_EXE	4005	Long non-transfer execution (i.e., format command)
LNXF_OUT	4006	Long non-transfer out
XFR_IN	4007	Transfer in
XFR_EXE	4008	Transfer execution (ie read or write)
XFR_OUT	4009	Transfer out
TAKEOVER_FAIL	5000	--
NO_RESP	5001	--
NO_RESP1	5001	Detected by CPU1
NO_RESP2	5002	Detected by CPU2
OS_FAIL	6000	Operating System Failure
SYSFAIL	7000	System Fatal Error
CBUF_PARITY	7001	Cache Buffer Detected Parity Error
CBUF_SERR	7002	--

Boot Messages

Boot messages can be extremely useful in troubleshooting situations. The following are examples of standard boot messages on Sun StorEdge T3+ arrays having no failures. Typical boot messages for the array workgroup and enterprise configurations appear below for reference.

This section consists of the following components:

Interpreting Boot Messages

Boot Message Acronyms

The acronyms used in boot messages are given in TABLE C-6.

TABLE C-6 Boot Message Acronyms

Files and Directories	Explanation
XPT	Refers to the SCSI transport driver
QLCF	Refers to the QLogic Fibre Channel driver
ISP2x00	The intelligent SCSI processor used in the T3
ECC	The error checking and correcting mechanism used in the Sun StorEdge T3+ array controller
XOR	The exclusive OR logic operation used in RAID 5
PGR	This is the persistent group reservation information user that exists when the Sun StorEdge T3+ array is attached to a cluster
ALPA	The arbitrated loop physical address assigned to each device on a FC-AL loop

Boot Message Bracket Placement

The synonymous boot message lines shown below gives the field service engineer information about how and where they are connected to the Sun StorEdge T3+ array:

```
Found units: [u1-ctr] u2-ctr
```

or

```
Found units: u1-ctr [u2-ctr]
```

The position of the brackets indicates which serial port is providing the output. Brackets around u1-ctr indicate that the boot messages are coming from the master controller's serial port. The field service engineer is using tip to connect to that controller's serial port. The same is true with u2-ctr.

Detecting FC-AL Ports and Reporting Firmware Status

Compare two sections of boot messages, one from the master controller and one from the alternate master controller:

```
Initializing loop 1 ISP2200 ... firmware status = 3
Detected 19 FC-AL ports on loop 1
Initializing loop 2 ISP2200 ... firmware status = 3
Detected 19 FC-AL ports on loop 2
```

and

```
Initializing loop 1 ISP2200 ... firmware status = 3
Detected 20 FC-AL ports on loop 1
Initializing loop 2 ISP2200 ... firmware status = 3
Detected 20 FC-AL ports on loop 2
```

These messages are generated by the ISP devices that service the back-end loops. They are polling the FC-AL loops for FC-AL devices. The first section of output from the master controller detects 19 FC-AL ports. The next section detects 20 FC-AL ports.

The missing port is actually the alternate master controller. It is missing because it has not completed its own boot process when the master controller polls FC-AL devices.

Once the alternate master boots, it also polls for FC-AL devices. Since the master controller and all the drives are already running, 20 (9 drives per Sun StorEdge T3+ array, a master controller, and an alternate master controller) devices are found on the FC-AL loop at this time.

The firmware status codes generated during the boot cycle can be good indicators of internally detected system and configuration problems. TABLE C-7 specifies the firmware status codes that can be reported through a serial port console during a Sun StorEdge T3+ array boot cycle:

TABLE C-7 Firmware Status Boot Messages

Status	Explanation
`firmware status = 0`	ISP is waiting for configuration process to complete
`firmware status = 1`	ISP is waiting for ALPA assignment
`firmware status = 2`	ISP is waiting for port login
`firmware status = 3`	ISP is ready and optimal
`firmware status = 4`	ISP has lost loop synchronization
`firmware status = 5`	ISP has experienced an unrecoverable error
`firmware status = 6`	ISP re-initialization
`firmware status = 7`	ISP is not participating in the loop

If the firmware status given in either of these boot messages is not 3, a drive or other component in the array could be faulty. The number of devices found is important when trying to determine the failing device. For example, if only half of the devices are found, a loop card or loop cable could be faulty.

The following message is generated by the ISP that services the front-end or host loop. A status of 7 (not participating) does not necessarily indicate a problem. The attached host might not be running, and thus cannot respond to the Sun StorEdge T3+ array.

```
Initializing host port u2p1 ISP2200 ... firmware status = 7
```

Sun StorEdge T3+ Array Workgroup Configuration

```
T3B-2
Starting POST
........
POST end

Starting...

T3B-EP Release 2.01 2002/07/30 16:33:52 (129.150.28.81)
Copyright (C) 1997-2002 Sun Microsystems, Inc.
All Rights Reserved.
```

```
Found units: [u1-ctr]
auto boot is enabled
hit the RETURN key within 3 seconds to cancel...

Starting...

T3B Release 2.00 2001/04/02 15:21:29 (129.150.28.81)
Copyright (C) 1997-2002 Sun Microsystems, Inc.
All Rights Reserved.

Initializing software...
Found units: [u1-ctr]
Default master is u1
Starting Heartbeats...
Assigning Select IDs: u1(1)
Initializing system drivers...
Initializing XPT component...
Initializing QLCF component...
Initializing loop 1 ISP2200 ... firmware status = 3
Detected 10 FC-AL ports on loop 1
Initializing loop 2 ISP2200 ... firmware status = 3
Detected 10 FC-AL ports on loop 2
Initializing SVD services...
Detected data cache size in system: 1GB
Testing ISP2200...  Passed
Testing ECC mechanism...  Passed
Testing XOR functions and datapaths... Passed
Cold Boot detected; destructive tests OK...
Testing data cache memory...  Passed
Initializing Cache Memory...
Initializing system DB structure...
Initializing configuration...
Initializing port configuration...
Initializing loop 2 to accept SCSI commands...
Mounting root volume...
Checking local file system...
Initializing network routes...
Read PGR data ... Done.
Starting Syslog Daemon...
System has 1 active controller(s)
Initializing TFTP...
Starting ftpd...
Starting telnetd...
Starting timed...
Starting pshd...
Starting httpd...
Starting snmpd...
Starting schd...
```

```
Checking disk positions...
Initializing host port u1p1 ISP2200 ... firmware status = 7
Host port u1p1 TARGET_ID = 0xffff (ALPA = 0x5)
Starting psh...


Login:
```

Sun StorEdge T3+ Array Enterprise Configuration

```
T3B-2
Starting POST
........
POST end



Starting...

T3B-EP Release 2.01 2002/07/30 16:33:52 (129.150.28.80)
Copyright (C) 1997-2002 Sun Microsystems, Inc.
All Rights Reserved.

Found units: [u1-ctr] u2-ctr
auto boot is enabled
hit the RETURN key within 3 seconds to cancel...



Starting...

T3B Release 2.01 2002/07/30 15:21:29 (129.150.28.80)
Copyright (C) 1997-2002 Sun Microsystems, Inc.
All Rights Reserved.



Initializing software...
Found units: [u1-ctr] u2-ctr
Default master is u1
Default alternate master is u2
Master coming up...
Starting Heartbeats...
Assigning Select IDs: u1(1) u2(2)
Initializing system drivers...
Initializing XPT component...
Initializing QLCF component...
Initializing loop 1 ISP2200 ... firmware status = 3
Detected 19 FC-AL ports on loop 1
Initializing loop 2 ISP2200 ... firmware status = 3
Detected 19 FC-AL ports on loop 2
Initializing SVD services...
Detected data cache size in system: 1GB
Testing ISP2200...  Passed
Testing ECC mechanism...  Passed
Testing XOR functions and datapaths... Passed
Cold Boot detected; destructive tests OK...
```

```
Testing data cache memory...  Passed
Initializing Cache Memory...
Initializing system DB structure...
Initializing configuration...
Initializing port configuration...
Initializing loop 2 to accept SCSI commands...
Mounting root volume...
Checking local file system...
Initializing network routes...
Read PGR data ... Done.
Starting Syslog Daemon...
Waiting for 1 slave controller(s) to come up...
u1: Configuring local data
u2: Initializing drives
System has 1 active controller(s)
Initializing TFTP...
Starting ftpd...
Starting telnetd...
Starting timed...
Starting pshd...
Starting httpd...
Starting snmpd...
Starting schd...
Checking disk positions...
Initializing host port u1p1 ISP2200 ... firmware status = 7
Host port u1p1 TARGET_ID = 0xffff (ALPA = 0x5)
Starting psh...


Login:
```

Sun StorEdge T3+ Array Enterprise Configuration as seen from the Alternate Master Controller

```
T3B-2
Starting POST
........
POST end


Starting...

T3B-EP Release 2.01 2002/07/30 16:33:52 (129.150.28.80)
Copyright (C) 1997-2002 Sun Microsystems, Inc.
All Rights Reserved.

Found units: u1-ctr [u2-ctr]
auto boot is enabled
hit the RETURN key within 3 seconds to cancel...


Starting...

T3B Release 2.01 2002/07/30 15:21:29 (129.150.28.80)
Copyright (C) 1997-2002 Sun Microsystems, Inc.
All Rights Reserved.

Initializing software...
Found units: u1-ctr [u2-ctr]
Default master is u1
Default alternate master is u2
Waiting for Master to come up...
Starting Heartbeats...
Initializing system drivers...
Initializing XPT component...
Initializing QLCF component...
Initializing loop 1 ISP2200 ... firmware status = 3
Detected 20 FC-AL ports on loop 1
Initializing loop 2 ISP2200 ... firmware status = 3
Detected 20 FC-AL ports on loop 2
Initializing SVD services...
Detected data cache size in system: 1GB
Testing ISP2200...  Passed
Testing ECC mechanism...  Passed
Testing XOR functions and datapaths... Passed
Cold Boot detected; destructive tests OK...
Testing data cache memory...  Passed
Initializing Cache Memory...
```

```
Initializing loop 2 to accept SCSI commands...
Starting Syslog Daemon...
Waiting for configuration data from master...
Initializing host port u2p1 ISP2200 ... firmware status = 7
Host port u2p1 TARGET_ID = 0xffff (ALPA = 0x5)
Starting psh...

Login:
```

Task List

Tasks on a Sun StorEdge T3+ array correspond to processes on a Solaris system. The following are typical Sun StorEdge T3+ array tasks.

TMRT: Timer Task - Handles fru removal time-out

LXR0: Handles incoming messages from loop card serial port, one for each loop card.

LXR1: Handles incoming messages from loop card serial port, one for each loop card.

HBTT: Heartbeat Task

LPCT: Loop card monitor task

CFGT: configuration task

WXFT: WriteTransferTask, waits for command set completion

SX01: StartTransferTask, Waits for the first command set to complete for the stripe and the head of the stripe order list

XFRT: Waits for a command, decompose it into stripes, and sets each stripe to the stripe requestor task

MXFT: Mirror transfer task

HS01: Simulates host I/Os to configured volumes

SMON: Handles events which effect cache mirroring

FCC0: ScsiPortCmdTask: Port task to handle host commands

FCC2: ScsiPortCmdTask: Backend loop mirror task

SIMT: Brings ISP back online (part of sim reset)

SVDT: Handles backend loop link events such as LIPs, loop up, loop down, etc.

SVHT: Handles front-end loop link events such as LIPs, loop up, loop down, etc

SDFT: Handles path and loop failover events

ONDG: Executes back-end loop diagnostics

TMON: Monitors disk temperature

IPCS: For multi-controller inter processor communication

IPCR: Partner to IPCS

LT00: Handles ‘long’ transfer command execution

LNXT: Handles ’long’ non-xfr command execution. These are commands that take a long time, like Reconstruct

MNXT: Handles ’medium’ non-transfer command execution

SNXT: Handles ’short’ non-transfer command execution

SCHD: Schedule manager

Ftpd: FTP daemon

ANNT: Wait for announce string and display is [syslog daemon]

TIME: Time Daemon

HT00: Process HTTP connections

HTPD: Listen for HTTP connections

SNMP: Process SNMP requests

Pshd: Shell Daemon - This spawns individual shell task

Pshc: Execute shell commands

Tnpd: Telnet daemon

Internal Sun StorEdge T3+ Array AL_PA / LID / LOOP Map

TABLE C-8 Internal Sun StorEdge T3+ Array AL_PA / LID / LOOP Map

Device	al_pa	loop_id	Target	LID	Order	Loop
u1d3	d5	0xa	10	3	1	0
u1d2	d6	0x9	9	2	2	0
u1d1	d9	0x8	6	1	3	0
u2d9	36	0x62	98	17	4	1
u2d8	c5	0x17	23	16	5	1
u2d7	c6	0x16	22	15	6	1
u2ctr	e8	0x01	1	N / A	7	N / A
u2d3	cb	0x12	18	11	8	0
u2d2	cc	0x11	17	10	9	0
u2d1	cd	0x10	16	9	10	0
u2d6	c7	0x15	21	14	11	1
u2d5	c9	0x14	20	13	12	1
u2d4	ca	0x13	19	12	13	1
u1d6	d2	0xd	13	6	14	1
u1d5	d3	0xc	12	5	15	1
u1d4	d4	0xb	11	4	16	1
u1d9	39	0x61	97	18	17	1
u1d8	ce	0xf	15	8	18	1
u1d7	d1	0xe	14	7	19	1
u1ctr	ef	0x00	0	N / A	20	N / A

SCSI Virtual Disk Driver (SVD) Error Definitions

TABLE C-9 SVD Disk Error Definitions

Opcode	Error
0x0	Request in progress
0x1	Completed without error
0x2	Retry attempted
0x3	Completed with error
0x4	Retries exhausted
0x5	LBA out of range
0x6	I/O enqueue failure
0x7	Invalid command specified
0x8	resource not available
0x9	Invalid command specified
0xA	Device already open
0xB	Device exclusively opened
0xC	Resource not available
0xD	On-disk label not found
0xE	Invalid path specified
0xF	Flush in progress
0x10	Device is not present
0x11	Device is not online
0x12	Command(s) active
0x13	Failover in progress
0x14	Device is broken
0x15	Device is unavailable

Stripe Type Messages

Stripe type messages report the I/O operation that was being performed when the stripe (RAID stripe) message occurred. These do not necessarily indicate an I/O operation failure. These messages are found in the `syslog` file.

TABLE C-10 Stripe Type Messages

Stripe Type Message	Description
	Control Stripe = 0x0100
0101	Dummy
0102	No_Alternate
0103	Alt_Possible
0104	Using_Alternate
	Header Stripe = 0x0200
0201	CCR_Header
0202	Asynch_Header
	Read Stripe = 0x0400
0401	RAID-1_Read
0402	RAID-1_Recon_Read
0403	RAID-1_Read_Check
0404	Data_Source
0405	RAID-5_Small_Read
0406	RAID-5_Recon_Read
0407	RAID-5_Rebuild_Read
0408	RAID-5_Stripe_Read
	Write Stripe = 0x0800
0801	Cache_Write
0802	Cache_Insert
	Write Disk Stripe = 0x0810
0811	RAID-0 Write to Disk
0812	RAID-1 Write to Disk
0813	RAID-0 Insert into disk Block

TABLE C-10 Stripe Type Messages *(Continued)*

Stripe Type Message	Description
0814	RAID-1 Insert into disk Block
0815	RAID-5 RMW to Disk
0816	RAID-5 Recon Write
0817	RAID-5 Stripe Write
081	RAID-5 RMW Insert
0819	RAID-5 Recon Insert
0820	Data Sink Recon Stripe = 0x1000
1001	Copy Recon (drv <> stdby)
1002	RAID-1 Recon
1003	RAID-0 Insert into disk Block Internal Stripe = 0x2000
2001	RAID-0 Write Data Init
2002	RAID-1 Write Data Init
2003	RAID-5 Write Data Init Flush Stripe = 0x2020
2021	RAID-0 Cache Flush
2022	RAID-1 Cache Flush
2023	RAID5 RMW Cache flush
2024	RAID-5 Recon Cache Flush
2025	RAID-5 Stripe flush Verify Stripe = 0x4000
4001	RAID-0 Read Verify
4002	RAID-1 Read Verify
4003	RAID-5 Read Verify

SCSI Command Set

A partial list of the SCSI commands available with the Sun StorEdge T3+ array are given in TABLE C-11. For a complete list of the commands see:

`http://www.t10.org/lists/op-num.htm)`

TABLE C-11 SCSI Command Set

Opcode	Commands	Supported	Notes
0x08	READ(6)	yes	
0x28	READ(10)	yes	
0x0A	WRITE(6)	yes	
0x2A	WRITE(10)	yes	
0x2E	WRITE AND VERIFY	yes	
0x2F	VERIFY	yes	
0x00	TEST UNIT READY	yes	
0x0	REQUEST SENSE	yes	
0x07	REASSIGN BLOCKS	no	
0x12	INQUIRY	yes	
0x16	RESERVE(6)	yes	
0x56	RESERVE(10)	yes	
0x17	RELEASE(6)	yes	
0x57	RELEASE(10)	yes	
0x1B	START STOP UNIT	yes	
0x25	READ CAPACITY	yes	
0x1D	SEND DIAGNOSTIC	yes	
0x1A	MODE SENSE(6)	yes	
0x5A	MODE SENSE(10)	yes	
0x15	MODE SELECT(6)	yes	
0x55	MODE SELECT(10	yes	
0xA0	REPORT LUNS	yes	
0x5E	PERSISTENT RESERVE IN	yes	

TABLE C-11 SCSI Command Set *(Continued)*

Opcode	Commands	Supported	Notes
0x5F	PERSISTENT RESERVE OUT	ye	
0xD0	LUN FAILOVER	yes	Vendor specific
0x3C	READ BUFFER	yes	Available in 1.18/2.0.1
0x3B	WRITE BUFFER	yes	Available in 1.18/2.0.1
0x4D	LOG SENSE	no	
0x4C	LOG SELECT	no	

Arbitrated Loop Physical Addresses (AL_PA) and Loop IDs

TABLE C-12 Arbitrated Loop Physical Addresses and Loop IDs

AL_PA (hex	SEL_ID (hex	Target (dec)	AL_PA (hex	SEL_ID (hex	Target (dec)	AL_PA (hex	SEL_ID (hex	Target (dec)
EF	00	0	A3	2B	43	4D	56	86
E8	01	1	9F	2C	44	4C	57	87
E4	02	2	9E	2D	45	4B	58	88
E2	03	3	9D	2E	46	4A	59	89
E1	04	4	9B	2F	47	49	5A	90
E0	05	5	98	30	48	47	5B	91
DC	06	6	97	31	49	46	5C	92
DA	07	7	90	32	50	45	5D	93
D9	08	8	8F	33	51	43	5E	94
D6	09	9	88	34	52	3C	5F	95
D5	0A	10	84	35	53	3A	60	96
D4	0B	11	82	36	54	39	61	97
D3	0C	12	81	37	55	36	62	98
D2	0D	13	80	38	56	35	63	99
D1	0E	14	7C	39	57	34	64	100
CE	0F	15	7A	3A	58	33	65	101
CD	10	16	79	3B	59	32	66	102
CC	11	17	76	3C	60	31	67	103
CB	12	18	75	3D	61	2E	68	104
CA	13	19	74	3E	62	2D	69	105
C9	14	20	73	3F	63	2C	6A	106
C7	15	21	72	40	64	2B	6B	107
C6	16	22	71	41	65	2A	6C	108
C5	17	23	6E	42	66	29	6D	109

TABLE C-12 Arbitrated Loop Physical Addresses and Loop IDs *(Continued)*

AL_PA (hex	SEL_ID (hex	Target (dec)	AL_PA (hex	SEL_ID (hex	Target (dec)	AL_PA (hex	SEL_ID (hex	Target (dec)
C3	18	24	6D	43	67	27	6E	110
BC	19	25	6C	44	68	26	6F	111
BA	1A	26	6B	45	69	25	70	112
B9	1B	27	6A	46	70	23	71	113
B6	1C	28	69	47	71	1F	72	114
B5	1D	29	67	48	72	1E	73	115
B4	1E	30	66	49	73	1D	74	116
B3	1F	31	65	4A	74	1B	75	117
B2	20	32	63	4B	75	18	76	118
B1	21	33	5C	4C	76	17	77	119
AE	22	34	5A	4D	77	10	78	120
AD	23	35	59	4E	78	0F	79	121
AC	24	36	56	4F	79	08	7A	122
AB	25	37	55	50	80	04	7B	123
AA	26	38	54	51	81	02	7C	124
A9	27	39	53	52	82	01	7D	125
A7	28	40	52	53	83	00	7E	126
A6	29	41	51	54	84	--	7F	127
A5	2A	42	4E	55	85	--	--	--

Note – The values are intentionally from lowest to highest priority. AL_PA = 00 is reserved for the FL_Port; -- is not available.

Source:

`ftp://ftp.t11.org/t11/member/fc/al/fcal44p.asc`

Calculating Port and Loop ids:

port_local = 3 x (encl_id - 1) + port_loop

To get the isp port# on a ctrl where:

encl_id = 1,2,...8

port_loop = 0,1,2

loop_id = (encl_id - 1)—This is the isp_id(->alpa) on each of the 2 back-end isp's on a ctrl (see chart at end of file)

Sense Key Explanations

Sense keys are returned from devices when issued a REQUEST SENSE command. They return more detailed information on a problem which occurred with a previous command. Here are the definitions of Sense keys as defined in the SCSI 2 proposed standard:

0xB ABORTED COMMAND

This indicates that the target aborted the command. The initiator may be able to recover by trying the command again.

0x8 BLANK CHECK

This indicates that a write-once device or a sequential-access device encountered blank medium or format-defined-end-of-data indication while reading or a write-once device encountered a non-blank medium while writing.

0xA COPY ABORTED

This indicates a COPY, COMPARE, or COPY AND VERIFY command was aborted due to an error condition on the source device, the destination device, or both.

0x7 DATA PROTECT

This indicates that a command that reads or writes the medium was attempted on a block that is protected from this operation. The read or write operation is not performed.

0xC EQUAL

This indicates a SEARCH DATA command has satisfied an equal comparison.

0x4 HARDWARE ERROR

This indicates that the target detected a non-recoverable hardware failure (for example, controller failure, device failure, parity error, etc.) while performing the command or during a self test.

0x5 ILLEGAL REQUEST

This indicates that there was an illegal parameter in the command descriptor block or in the additional parameters supplied as data for some commands (FORMAT UNIT, SEARCH DATA, etc.). If the target detects an invalid parameter in the command descriptor block, then it shall terminate the command without altering the medium. If the target detects an invalid parameter in the additional parameters supplied as data, then the target may have already altered the medium. This sense key may also indicate that an invalid IDENTIFY message was received.

0x3 MEDIUM ERROR

This indicates that the command terminated with a non-recovered error condition that was probably caused by a flaw in the medium or an error in the recorded data. This sense key may also be returned if the target is unable to distinguish between a flaw in the medium and a specific hardware failure (sense key 0x4).

0xE MISCOMPARE

This indicates that the source data did not match the data read from the medium.

0x0 NO SENSE

This indicates that there is no specific sense key information to be reported for the designated logical unit. This would be the case for a successful command or a command that received CHECK CONDITION or COMMAND TERMINATED status because one of the file mark, EOM or ILI bits is set to one.

0x2 NOT READY

This indicates that the logical unit addressed cannot be accessed. Operator intervention may be required to correct this condition.

0x1 RECOVERED ERROR

This indicates that the last command completed successfully with some recovery action performed by the target. Details may be determinable by examining the additional sense bytes and the information field. When multiple recovered errors occur during one command, the choice of which error to report (first, last, most severe, etc) is device specific.

0xF RESERVED

0x6 UNIT ATTENTION

This indicates that the removable medium may have been changed or the target has been reset.

0x9 VENDOR-SPECIFIC

This sense key is available for reporting vendor specific conditions.

0x0 VOLUME OVERFLOW

This indicates that a buffered peripheral device has reached the end-of-partition and data may remain in the buffer that has not been written to the medium. A RECOVER BUFFERED DATA command(s) may be issued to read the unwritten data from the buffer.

APPENDIX **D**

Sun StorEdge T3+Array System Commands

This appendix lists the commands supported by the Sun StorEdge T3+ array and is divided into the following sections:

Commands List

To view the available command-line interface (CLI) commands on the array, type `help` at the prompt.

```
:/:<184>help
arp     cat     cd      cmp     cp      date    echo    head
help    ls      mkdir   mv      ping    pwd     rm      rmdir
tail    touch
boot    disable disk    enable  fru     id      logger  lpc
more    ntp     passwd  port    proc    reset   set     shutdown
sync    sys     tzset   ver     vol     volslice        ep
refresh route   ofdg    lun     hwwn
```

Note – Use the login prompt to set the IP address, netmask, and hostname instead of using the EP prompt. Setting these parameters at the EPROM level will be lost.

To display command syntax, use the *command-name* `help` command. For example, for information on the `reset` command, type:

```
:/:<1> reset help
reset  - reset system (reentrant, not locked)
```

TABLE D-1 contains an alphabetical listing of the CLI commands supported by the array.

See the *Sun StorEdge T3+ Array Administrator's Manual* for a detailed description of each command's syntax, options, and arguments.

TABLE D-1 Commands Listed in Alphabetical Order

Command	Description
`boot`	Boot system
`disable`	Disable certain FRUs
`disk`	Disk administration
`enable`	Enable certain FRUs
`ep`	Program the flash EPROM
`fru`	Display the FRU information
`help`	Display reference manual pages
`id`	Display FRU identification summary
`logger`	Generate messages to the `syslog` in the unit; dump the reset log and display system crash information
`lpc`	Get interconnect card property
`port`	Configure the interface port
`proc`	Displays status of outstanding `vol` processes
`refresh`	Start/stop battery refreshing or display its status
`reset`	Reset system
`set`	Display or modify the set information
`shutdown`	Shut down array or partner group
`sys`	Display or modify the system information
`ver`	Display software version
`vol`	Display or modify the volume information

FRU Identifiers

Many commands use a FRU identifier to refer to a particular FRU in an array. This identifier contains a unit constant (`u`), the unit number (e*ncid*), the FRU constant (`ctr` for controller card, `pcu` for power and cooling unit, `l` for interconnect card, `d` for disk drive) and the FRU number (*n*). TABLE D-2 lists the possible FRU variables as they appear in this appendix

TABLE D-2 FRU Identifiers

FRU	Identifier	Unit number
Controller card	u*encid*`ctr`	*encid* = unit number (1, 2, ...)
Power and cooling unit	u*encid*`pcu[1|2]`	*encid* = unit number (1, 2, ...) *n* = pcu number (1, 2)
Interconnect card	u*encid*`l[1|2]`	*encid* = unit number (1, 2, ...) *n* = interconnect card number (1, 2)
Disk drive	u*encid*`d`*n*	*encid* = unit number (1, 2, ...) *n* = disk drive number (1 to 9)

APPENDIX **E**

FC-AL Loop Identifiers

This Appendix lists the FC-AL loop identifies by AL_PA (hex), Switch (hex) and Setting (decimal) values.

The values are listed from lowest to highest priority. The AL_PA value of 00 is reserved for an FL_PORT. The value -- is not available.

TABLE E-1 Assigned Loop Identifier

AL_PA (hex)	Switch (hex)	Setting (dec)	AL_PA (hex)	Switch (hex)	Setting (dec)	AL_PA (hex)	Switch (hex)	Setting (dec)
EF	00	2	9F	2C	44	4B	58	88
E8	01	1	9E	2D	45	4A	59	89
E4	02	2	9D	2E	46	49	5A	90
E2	03	3	9B	2F	47	47	5B	91
E1	04	4	98	30	48	46	5C	92
E0	05	5	97	31	49	45	5D	93
DC	06	6	90	32	50	43	5E	94
DA	07	7	8F	33	51	3C	5F	95
D9	08	8	88	34	52	3A	60	96
D6	09	9	84	35	53	39	61	97
D5	0A	10	82	36	54	36	62	98
D4	0B	11	81	37	55	35	63	99
D3	0C	12	80	38	56	34	64	100
D2	0D	13	7C	39	57	33	65	101
D1	0E	14	7A	3A	58	32	66	102
CE	0F	15	79	3B	59	31	67	103

TABLE E-1 Assigned Loop Identifier

AL_PA (hex)	Switch (hex)	Setting (dec)	AL_PA (hex)	Switch (hex)	Setting (dec)	AL_PA (hex)	Switch (hex)	Setting (dec)
CD	10	16	76	3C	60	2E	68	104
CC	11	17	75	3D	61	2D	69	105
CB	12	18	74	3E	62	2C	6A	106
CA	13	19	73	3F	63	2B	6B	107
C9	14	20	72	40	64	2A	6C	108
C7	15	21	71	41	65	29	6D	109
C6	16	22	6E	42	66	27	6E	110
C5	17	23	6D	43	67	26	6F	111
C3	18	24	6C	44	68	25	70	112
BC	19	25	6B	45	69	23	71	113
BA	1A	26	6A	46	70	1F	72	114
B9	1B	27	69	47	71	1E	73	115
B6	1C	28	67	48	72	1D	74	116
B5	1D	29	66	49	73	1B	75	117
B4	1E	30	65	4A	74	18	76	118
B3	1F	31	63	4B	75	17	77	119
B2	20	32	5C	4C	76	10	78	120
B1	21	33	5A	4D	77	0F	79	121
AE	22	34	59	4E	78	08	7A	122
AD	23	35	56	4F	79	04	7B	123
AC	24	36	55	50	80	02	7C	124
AB	25	37	54	51	81	01	7D	125
AA	26	38	53	52	82			
A9	27	39	52	53	83	00	7E	126
A7	28	40	51	54	84	--	7F	127
A6	29	41	4E	55	85			
A5	2A	42	4D	56	86			
A3	2B	43	4C	57	87			

APPENDIX **F**

Sun StorEdge T3+ Array Configuration Worksheets

This chapter contains a blank worksheet for the qualified service provider to make notes at each customer site and contains the following sections:

Worksheets

The following information is required to successfully troubleshoot a Sun StorEdge T3+ array.

Use this worksheet to access the data, Ethernet, and TFTP connections from the application, management, and TFTP host system(s). The application, management, and TFTP host can all be resident on the same server.

Supervisor access is required for all hosts during troubleshooting.

Host types are defined as the following:

Application host	The application host utilizes the FC-AL connection as a data path to and from the Sun StorEdge T3+ array.
Management host	The management host administers configuration and health monitoring of the Sun StorEdge T3+ array, through a network connection.
TFTP host	The TFTP host is used to download bootcode to the Sun StorEdge T3+ array, through a network connection.

System Information Worksheets

The following information should be documented before troubleshooting any Sun StorEdge T3+ array. Make copies of this blank form, and complete it for each Sun StorEdge T3+ array.

TABLE F-1 Sun StorEdge T3+ array Information Worksheet

	Management Host	Application Host	TFTP Host
Host ID			
Host name			
Host IP address			
Gateway IP address			
Sun StorEdge T3+ IP address			
Sun StorEdge T3+ array name			
TFTP host address			
OS/patch revision level			
VERITAS DMP release			
Primary application			
Sun Storage Automated Diagnostic Environment release			

TABLE F-1 Sun StorEdge T3+ array Information Worksheet *(Continued)*

	Management Host	Application Host	TFTP Host
	Legend:		
	Required Field		
	Optional Field		
	Not Applicable		

Glossary

A

administrative domain Partner groups (interconnected controller units) that share common administration through a master controller.

alternate master unit The secondary array unit in a partner group that provides failover capability from the master unit.

alternate pathing (AP) A mechanism that reroutes data to the other array controller in a partner group upon failure in the host data path. Alternate pathing requires special software to perform this function.

auto cache mode The default cache mode for the Sun StorEdge T3+ array. In a fully redundant configuration, cache is set to write-behind mode. In a nonredundant configuration, cache is set to write-through. Read caching is always performed.

auto disable The Sun StorEdge T3+ array default that automatically disables a disk drive that has failed.

auto reconstruction The Sun StorEdge T3+ array default that automatically reconstructs data onto a new disk drive from one of the other drives.

B

buffering The process of data transfer between the host and the drives.

C

command-line interface (CLI) The interface between the Sun StorEdge T3+ array's pSOS operating system and the user in which the user types commands to administer the array.

controller unit A Sun StorEdge T3+ array that includes a controller card. It can be use as a standalone unit or configured with other Sun StorEdge T3+ arrays.

E

erasable programmable read-only memory (EPROM) Memory stored on the controller card; useful for stable storage for long periods without electricity while still allowing reprogramming.

F

Fibre Channel-Arbitrated Loop (FC-AL) A 100 MB/sec serial channel, which allows connection of multiple devices (disk drives and controllers).

field-replaceable unit (FRU) A component that is easily removed and replaced by a field service engineer or a system administrator.

G

gigabit interface converter (GBIC) An adapter used on an SBus card to convert fiber-optic signal to copper.

gigabyte (GB or Gbyte) One gigabyte is equal to one billion bytes (1x109).

graphical user interface (GUI) A software interface that enables configuration and administration of the Sun StorEdge T3+ array using a graphic application.

H

hot spare A drive in a RAID 1 or RAID 5 configuration that contains no data and acts as a standby in case another drive fails.

hot-swap The characteristic of a field-replaceable unit (FRU) to be removed and replaced while the system remains powered on and operational.

I

input/output operations per second (IOPS) A performance measurement of the transaction rate.

interconnect cable An FC-AL cable with a unique switched loop architecture that is used to interconnect multiple Sun StorEdge T3+ arrays. Sometimes referred to as a *loop cable*.

interconnect card A array component that contains the interface circuitry and two connectors for interconnecting multiple Sun StorEdge T3+ array units. Sometimes referred to as a *loop card*.

L

light emitting diode (LED) A device that converts electrical energy into light that is used to display activity.

logical unit number (LUN) One or more drives that can be grouped into a unit; also called a *volume*.

loop cable Interconnect cable.

loop card Interconnect card.

M

master unit The main controller unit in a partner group configuration.

media access control (MAC) address A unique address that identifies a storage location or a device.

media interface adapter (MIA) An adapter that converts fiber-optic light signals to copper.

megabyte (MB or Mbyte) One megabyte is equal to one million bytes (1x106).

megabytes per second (MB/sec) A performance measurement of the sustained data transfer rate.

P

parity Additional information stored with data on a disk that enables the controller to rebuild data after a drive failure.

partner group A pair of interconnected controller units.

power and cooling unit A FRU component in the Sun StorEdge T3+ array. The unit contains a power supply, cooling fans, and an integrated UPS battery. A Sun StorEdge T3+ array contains two power and cooling units.

pSOS A real-time operating system used as the primary operating system for the Sun StorEdge T3+ array.

R

read caching Data for future retrieval, to reduce disk I/O as much as possible.

reliability, availability, serviceability (RAS) Product features that include high availability, easily serviced components that are very dependable.

redundant array of independent disks (RAID) A configuration in which multiple drives are combined into a single virtual drive, to improve performance and reliability.

S

Simple Network Management Protocol (SNMP) A protocol for remotely managing a computer network.

synchronous dynamic random access memory (SDRAM) A form of dynamic random access memory (DRAM) that can run at higher clock speeds than conventional DRAM.

system area Located on the disk drive label, the space that contains configuration data, boot firmware, and file system information.

U

uninterruptable power source (UPS) A component within the power and cooling unit. The UPS supplies power from a battery in the case of an AC power failure.

unit interconnect card (UIC) *See* Interconnect Card.

V

volume Also called a logical unit number (LUN), a volume is one or more drives that can be grouped into a unit for data storage.

W

write caching Data used to build up stripes of data, eliminating the read-modify-write overhead. Write caching improves performance for applications that are writing to disk.

Index

D

E

F

G

H

I

L

M

N

O

P

R

S

T

U

V

W