



# The Sun StorEdge™ T3 Array: Installation, Configuration, and Monitoring Best Practices

---

*By Ted Gregg - Customer Problem Resolution  
Engineering*

*Sun BluePrints™ OnLine - October 2001*



<http://www.sun.com/blueprints>

**Sun Microsystems, Inc.**  
901 San Antonio Road  
Palo Alto, CA 94303 USA  
650 960-1300 fax 650 969-9131

Part No.: 816-1937-10  
Revision 01, 09/19/01  
Edition: October 2001

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303 U.S.A. All rights reserved.

This product or document is protected by copyright and 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, Sun BluePrints, Sun StorEdge, StorTools, Sun Enterprise, SunVTS, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

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.

**RESTRICTED RIGHTS:** Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

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 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et 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, Sun BluePrints, Sun StorEdge, StorTools, Sun Enterprise, SunVTS, 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.

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.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please  
Recycle



Adobe PostScript

# The Sun StorEdge™ T3 Array: Installation, Configuration, and Monitoring Best Practices

---

The StorEdge™ T3 Array represents the third generation of Fibre Channel disk arrays from Sun Microsystems. Based on a modular brick architecture, the Sun StorEdge T3 Array incorporates multiple hardware and software features to maximize reliability, availability, and serviceability (RAS). To fully realize the benefits of the capabilities built into the array, it must be installed, configured, and monitored with best practices for RAS considered. This article discusses these best practices and includes both the array configuration recommendations and host system configuration recommendations. Brief descriptions are included for some of the available software installation and monitoring tools— StorTools™ 3.3 and Sun StorEdge™ Component Manager 2.2 software.

Requirements for eliminating single-points-of-failure (SPOFs) from storage solutions that employ the arrays are provided. Recommended array hardware RAID and host-based software RAID configurations are also discussed. The reader is encouraged to review the documents referenced in this article before attempting to apply the information presented.

Although the array is sold and supported by multiple vendors on a variety of operating systems, recommendations made here only apply to the Solaris™ Operating Environment (Solaris OE). This article only addresses single-host configurations using direct-attached or FC-AL hub-attached arrays with SBus or PCI FC100 host adapters.

---

## Sun StorEdge T3 Array Overview

The Sun StorEdge T3 Array is a high performance, modular, scalable hardware RAID storage device containing nine disk drives and an intelligent controller. It is sold by Sun in two basic configurations: a single controller unit “workgroup” model—the T3WG (see FIGURE 1), and a dual controller unit “enterprise” model—the T3ES (see FIGURE 2). The dual controller T3ES is fundamentally two T3WGs cabled together to act as one storage device. The T3ES is also referred to as a “partner group,” because it is two controller units acting as one.

In partner group configurations, one controller unit is referred to as the “master,” while the other controller unit is called the “alternate master.” Both units function essentially the same for application host I/O requests. The master unit has the additional responsibility of providing the TCP services included with the array and acts as the management interface for the pair. The T3ES offers the advantage of cache mirroring and multi-path support over the T3WG, eliminating the controller as a SPOF to provide a fully redundant storage solution.



FIGURE 1 T3WG (Single Brick)




FIGURE 2 T3ES (Partner Group)

Both T3WG and T3ES configurations support the ability to create denser storage solutions by using FC-AL hubs. The disadvantage of using hubs over direct attach, however, is performance and reliability decrease, while troubleshooting complexity and mean-time-to-repair (MTTR) increase.

Each array controller unit consists of a chassis with a passive midplane that houses 4 types of Field Replaceable Units (FRUs): 2 loop cards (unit interconnect cards), 2 power/cooling units, 9 disk drives, and 1 controller card. Arrays are not sold without all FRUs in place. This is done to ensure redundancy of components and proper airflow for cooling. If a FRU failure is reported, the FRU must not be removed until a replacement is available. The array operating system allows any of these FRUs to be removed for up to 30 minutes for serviceability. After 30 minutes, the software initiates an array shutdown to prevent overheating. Since units cabled together act as a single managed system, the shutdown could affect multiple disk trays, so adherence to maintenance procedures is extremely important.

The array includes out-of-band management, monitoring, and SNMP capability via standard Ethernet. This means the arrays can be configured and monitored from systems on a network separate from the systems using the arrays as a storage device. They can also be monitored using the same applications used to monitor the rest of the enterprise—HP OpenView, BMC Patrol, IBM Tivoli, and CA Unicenter, to name a few. This will become more important as the principles of direct attached storage are replaced by those of Storage Area Networks (SANs).

---

## Installation

The Sun StorEdge T3 Array installation process is covered in depth in the references provided at the end of this article. This section covers the highlights and makes specific recommendations that may not be included in the installation manuals. This article is by no means a substitute for the published documentation, and every effort should be made to comply with the procedures provided therein.

## Planning and Host Definitions

All good installations start with a comprehensive plan that identifies everything necessary to support the installation effort. First, the host systems required to support the arrays must be understood. The array host systems fall into three categories: management hosts, application hosts, and TFTP hosts.

- **Management Host:** A management host is used to administer, configure, and monitor the arrays via the arrays's 10BaseT Ethernet port.

- **Application Host:** An application host uses the array fibre channel connection to read and write data stored on the array. It is also referred to as a data host.
- **TFTP Host:** A TFTP host is a system with the array controller firmware (boot code) that can be used to boot an array. This is also accomplished via the array Ethernet connection. This host is usually only required for troubleshooting purposes.

There is no requirement that the three logical hosts defined here be different physical systems; one system can be used for all three purposes. For large configurations, the functions should be separated so that the application hosts are not unnecessarily burdened with administrative tasks and monitoring overhead.

Once the management, application, and TFTP hosts have been identified, array IP addresses and hostnames must be assigned. Although a partner group requires only one IP address, it is good practice to reserve an IP for each controller unit, should TFTP booting or partner group splitting become necessary. Likewise, it is good practice to record the MAC address (which can be found on the white pull-out tab to the left of disk 1 on the array chassis) of each unit. The MAC address is necessary to assign the array's initial IP address using Reverse Address Resolution Protocol (RARP), which will be discussed later.

## Array Mounting and Cabling

The Sun StorEdge T3 Array can be physically installed two ways: on a table top or in a rack. A table top installation simply requires a surface sturdy enough to hold the weight of each 67 pound array placed on it. The arrays can only be stacked two arrays high to prevent damage to the bottom array. They must be placed rubber foot side down; they should not be placed on their sides. If a partner group is being table top mounted, it should be cabled in the default manner, with the master unit on the bottom as shown in Figure 2. The master unit in a group of controller units is the one whose loop cards have connections only to the right "two-dot" port of each loop card. Attempting to place two units side-by-side and cabling them as a partner group could damage the loop cables.

Partner groups cabled in racks should follow the same convention. Master units should be on the bottom. If the arrays are being mounted in a standard Sun StorEdge 72" rack, requirements include the addition of a fan tray and an EMI front screen door. Since the arrays are front-to-back cooled, there is no need to keep a 2" gap between the cabinets' left side vent and its adjacent vertical surface, unless a side-cooled device, like a Sun Enterprise™ 4500 server, is mounted in the rack with the arrays. Refer to the *Rackmount Placement Matrix* and *Sun StorEdge Expansion Cabinet* manuals listed in the reference material to make sure everything is rack mounted properly.

One of the key availability features of the array is redundant, hot-swappable power supplies with integrated fans and batteries called power/cooling units (PCUs). For optimal availability, they should be powered from separate grids or circuits to prevent unnecessary outages due to power loss. In a Sun StorEdge 72" rack, PCU1 from each array should be powered by one sequencer, and PCU2 by the other. Whether wired on-site or pre-configured at the factory, they should be tested to make sure they are connected properly.

Each array controller has a copper FC-AL port, 10BaseT Ethernet port, and a serial port. The copper FC-AL port should be connected to the included Media Interface Adapter (MIA), which is, in turn, connected to fiber cabling and routed to the data host, possibly through a hub or switch. Data host connections should ensure that redundant paths are split between host adapters, system buses, I/O boards, and system boards to maximize performance and availability. Each Ethernet port should be connected to a common network which will be used to configure and administer the arrays. In a partner group only one port will be in use at any given time; however, both ports should be connected in the event a controller failover occurs.

Normally, the serial port is reserved for use by trained Sun service personnel. In fact, the *Sun StorEdge T3 Disk Tray Release Notes* recommend leaving it disconnected except during troubleshooting, or citing EMI and unauthorized access concerns. However, some enterprise customers with large installations may wish to connect the serial ports to a terminal server to facilitate monitoring and troubleshooting. If properly shielded racks are used and access is controlled, this facility can provide an invaluable tool for remote troubleshooting. In all cases, cleanly route and clearly label each fiber, Ethernet, or serial cable on both ends.

## Management Host Setup

The management host can be any workstation with access to the array administrative network. It is important to remember that all communication between the array and any hosts on its administrative network is not secure. The array has no method to encrypt any of its network services, and all passwords are transmitted in clear text. Although it is useful to include data hosts as well as management hosts on the array management subnet for logging purposes, security concerns may preclude this. The management host and the array network should be maintained separately from public networks wherever possible.

Initial configuration of the arrays, including assigning an IP address via RARP, should be done via the management host. If unable to assign an IP via RARP, the serial connection must be used to assign an IP. Refer to the *Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual* for detailed instructions on configuring RARP. After an IP has been assigned, `telnet` to the array and login as `root`. By default the password is not set.

Use the `set` command to view and change the array network parameters:

```
T3:/:<1>set
bootmode auto
bootdelay 3
sn XXXXXX
ip 10.0.0.1
netmask 255.255.255.0
gateway 0.0.0.0
tftp host 0.0.0.0
tftp file
hostname T3
vendor SLR-MI
model 375-0084-02-J-J9
revision 0210
logto *
loglevel 3
rarp on
mac XX:XX:XX:XX:XX:XX
```

If the IP was not set properly, it can be changed using the `set` command. Properly setting the IP, netmask, and gateway will allow the array to communicate on the network:

```
T3:/:<2>set ip <address>
T3:/:<3>set netmask <mask>
T3:/:<4>set gateway <address>
```

Changes to any of these three parameters requires an array reset to take effect:

```
T3:/:<5>reset -y
```

After network access has been established to an array, setting its root password is important, not only for security reasons, but to facilitate FTP. Without a root password, the array will not allow FTP access.

```
T3:/:<1>passwd
OLD password: <enter old password>
NEW password: <enter new password>
NEW password (confirm): <enter new password>
```



It is also important to set the time zone and date appropriately, so messages are logged properly and scheduled events occur as expected. If the time stamps on the array error messages do not correlate with those on the management and data hosts, troubleshooting can be much more difficult. The time zone and date are set using the `tzset` and `date` commands, respectively. A script called “t3time” that will synchronize an array’s time zone, date, and time with a host is available from your local Sun service representative.

The primary management interface, other than command line, is Component Manager, which should be installed on the management host. Component Manager is supported on Solaris OE and Microsoft Windows. Refer to the Component Manager references listed for installation details and patch requirements. Using Component Manager to monitor the arrays will be discussed later. Any SNMP management software to be used with the arrays should also be installed on the management host. Refer to the manufacturer's installation instructions to properly configure the software.

## Application Host Setup

The array’s application host should have the desired version of Solaris OE (2.6 or later) installed. The array related patches should be verified current. This can be done using the *Sun StorEdge T3 Disk Tray Release Notes* and SunSolve. Keeping host system patches and the array firmware up-to-date via a controlled process is essential to the overall availability of the system. The array contains four types of firmware: EProm, Controller (boot code), Loop Card, and Disk. Complete instructions for maintaining the array firmware are in the *Sun StorEdge T3 Disk Tray Administrator’s Guide*. The desired logical volume manager and StorTools must also be installed and patched. Logical volume manager configuration is discussed in the “Configuration” section of this article. If possible, application hosts should have access to the array management subnet to facilitate logging.

An important part of the installation process is equipment burn-in. Although arrays undergo significant factory testing, the packing and shipping process can contribute to early life component failure. To minimize this, all new installations should include a burn-in period of at least 24 hours. The array and host system burn-in is normally done using Sun's Validation and Test Suite (SunVTS™) and StorTools software. Complete instructions for using these tools can be found in the *SunVTS User’s Guide* and the *Sun StorEdge StorTools User’s Guide*. See the references section for more information.

## TFTP Host Setup

The TFTP host can be any system that supports TFTP with access to the array network. It is typically the same physical system as the management host. To set up a Solaris OE based TFTP host, create a `/tftpboot` directory, copy the array controller firmware into `/tftpboot`, edit `/etc/inetd.conf` to run `tftp`, and restart `inetd`. On the array, set the `tftphost` and `tftpfile` variables appropriately:

```
T3:/:<1>set tftphost <IP address>
T3:/:<2>set tftpfile <filename>
```

To actually boot an array using TFTP in a troubleshooting or testing scenario, change the bootmode to `tftp` and perform a reset:

```
T3:/:<3>set bootmode tftp
T3:/:<4>reset -y
```

It is important to note the differences between `auto` bootmode and `tftp` bootmode. When booting in `auto` mode, the array uses boot code written to the array disks' system area and plumbs the Ethernet interface of the master controller near the end of the boot cycle. In a partner group configuration, the alternate master controller's Ethernet interface is not activated, and only gets plumbed if the master controller fails or is manually disabled. In this failover scenario, the alternate master assumes the MAC address and IP address of the master and spawns the appropriate TCP services.

In `tftp` bootmode, both the master and alternate master controller's Ethernet ports are plumbed early in the boot cycle. This is necessary since the boot code has to be loaded over the network from the TFTP boot server. Thus, in this bootmode, each controller requires a separate IP address and will use its own MAC address. Failure to understand this can result in IP address conflicts and array boot failure when switching between `auto` and `tftp` bootmodes on partner groups. Since this bootmode is generally only used during troubleshooting, this can severely impact recovery efforts.

If unable to boot in `tftp` mode, check for IP address conflicts and stale ARP cache data on the TFTP boot server, and verify proper setup of the TFTP boot server and the array. For detailed instructions refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

---

# Configuration

Sun StorEdge T3 Array configurations can be lumped into two main categories—those *with* host-based mirroring, and those *without* host-based mirroring. Configurations relying on host-based mirroring should use array single bricks. Configurations where host-based mirroring is not employed should use array partner groups. This policy generally provides the best mix of RAS versus performance across a wide range of applications. This section provides insight for the reasons why. This section also makes recommendations for the appropriate selection of RAID levels and the proper setting of tunable array system parameters.

## Sun StorEdge T3 Array System Parameters

The Sun StorEdge T3 Array has been designed to make configuration straightforward and simple. The array caching algorithms adjust behavior based on the workload to simplify tuning. The base array configuration parameters are displayed and set using the `sys` command typed at the array command prompt:

```
T3:/:<1>sys
usage:  sys list
        sys stat
        sys blocksize <16k | 32k | 64k>
        sys cache <auto | writebehind | writethrough | off>
        sys mirror <auto | off>
        sys mp_support <none | rw | mpzio | std>
        sys rd_ahead <on | off>
        sys recon_rate <high | med | low>
```

The defaults values for these parameters are:

```
T3:/:<2>sys list
blocksize      : 64k
cache          : auto
mirror         : auto
mp_support     : none
rd_ahead      : on
recon_rate    : med
sys memsize   : 32 MBytes
cache memsize : 256 MBytes
```

For a complete discussion of these parameters, see the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

While the `sys` and `cache memsize` values are fixed, the other parameters are user configurable. The `blocksize` parameter determines the disk block layout, cache segment size, and the number of cache buffers available. It is global for all the array volumes (LUNs) configured in an array group. Changing this value requires existing LUNs be removed, thus it is important to set this value appropriately based on application type and I/O size before creating any array LUNs. Some general guidelines for selecting the appropriate blocksize are given in the following table:

<b>Application Type</b>	<b>Block Size</b>
Online Transaction Processing (OLTP) Internet service provider (ISP) Enterprise Resource Planning (ERP)	16K
NFS file system, version 2 Attribute-intensive NFS file system, version 3	32K
Data-intensive NFS file system, version 3 Decision Support Systems (DSS) Data Warehouse (DW) High Performance Computing (HPC)	64K

When selecting a blocksize, it is important to understand how drive size affects performance. An intensive OLTP application with small random I/O may perform satisfactorily with a 16 K blocksize on an array with 18 GB drives. Selecting the same parameters on an array with larger drives may yield performance that is unacceptable. In addition, blocksize affects reconstruction time. The smaller the blocksize, the longer the reconstruction time. Systems must be architected to ensure applications can withstand the performance impact of common FRU failures, such as disk drives.

The `rd_ahead` parameter controls disk read ahead into cache. It should also be set based on application I/O characteristics (i.e., `on` for a DSS database with sequential reads; `off` for an OLTP database with random reads). Unlike `blocksize`, `rd_ahead` can be changed on the fly, thus some simple experimentation with the user application should help determine the best setting. Cache read ahead is only

invoked if two consecutive application I/O blocks are requested by the host, at which time the remainder of the appropriate 64 KB array block is read into cache. For application I/O block sizes greater than or equal to 32KB, the `rd_ahead` parameter has no effect.

The `cache` and `mirror` parameters determine how the cache behaves under normal conditions and while operating with failed FRUs. Each controller has 256 MB of RAM that is used for caching read and write data. In cache mode `off`, no read or write caching is performed. In cache mode `writethrough`, only read caching is done. In cache mode `writebehind`, reads and writes are always cached, regardless of FRU status (Note: Disk drive failures will force the cache to `writethrough` mode due to bug 4410516). In cache mode `auto`, reads are always cached, but writes are only cached when FRU status is optimal.

Cache mirroring is a feature that causes write data to be mirrored on both controllers in a partner group to provide redundancy. It can only be set to `auto` if cache is set to `auto`. Both parameters should normally be set to `auto` to maximize availability. Specific instances where this is not recommended will be examined later.

The `mp_support` variable determines whether the array will export multiple paths to the data host(s). It should be set to `rw` whenever host-based multi-pathing software will be used. Obviously for Sun StorEdge T3 Array single brick configurations, this setting should be left at the default of `none`, since there is only one controller, therefore, only one path. Current supported multi-pathing software includes Dynamic Multi-Pathing (DMP) in Veritas Volume Manager version 3.0.2 and later, and Sun Alternate Pathing (AP) version 2.3.1 and later. With the release of array controller firmware 1.17, two new settings, `mpxio` and `std`, were added to the `mp_support` variable to support explicit LUN failover (ELF). ELF is a feature that allows future host applications to eliminate unintentional disk tray LUN failovers in multi-hosted configurations. These features are beyond the scope of this article.

The `recon_rate` parameter adjusts how much time is spent by the array doing data host I/O versus performing volume reconstruction following a disk failure. A setting of `high` will allow disk reconstruction to complete sooner, reducing exposure to an additional failure. However, it will also reduce the I/O rate and performance of the data host, potentially to an unacceptable level. Setting a value of `low` has the opposite effect. Unless there is good reason (i.e. proven benefits demonstrated by testing failures with the user application under normal workloads), it should be left at the default value.

## Sun StorEdge T3 Array RAID Levels

The Sun StorEdge T3 Array controller supports three hardware RAID levels: 0, 5, and 1 (mirroring). Up to two LUNs per array are configured as one of these three types. Refer to the *Sun StorEdge T3 Array Architecture* white paper for a good explanation of LUN configuration rules.

- RAID0: LUNs configured as RAID0 have data striped across 2 to 9 disks within an array. There is no protection against disk failure and no standby (hot spare) disk can be configured. This mode is normally only used in combination with data host-based mirroring.
- RAID5: LUNs configured as RAID5 have data striped across 3 to 9 disks within an array, with the option of using one disk as a standby. One segment of each stripe is dedicated to parity, thus a single disk failure will not result in data loss.
- RAID1: LUNs configured as RAID1 have data striped and mirrored across 2 to 9 drives within an array, with the option of using one disk as a standby. A single disk failure will not result in data loss. A double disk failure should not result in data loss, as long as two adjacent disks do not fail; however, the LUN may still become temporarily inaccessible.

Unless data host-based mirroring is to be employed in a configuration, RAID0 LUNs should not be used. For most applications, RAID5 performs well and provides an adequate level of protection. The advantage of RAID1 over RAID5 is approximately 20% better performance for small random writes. Since the data in a RAID1 LUN is mirrored within the same array, there is only a small increase in availability using RAID1 over RAID5. The primary advantage of RAID5 over RAID1 is more usable capacity within the same footprint, and thus lower cost.

The first choice in designing an array storage solution is deciding what RAID level is appropriate for the application. If RAID5 will yield satisfactory performance, then hardware RAID5 should be used. While it is technically possible to use data host-based software RAID5, the supportability, complexity, performance implications, and real world application usefulness preclude its discussion here. If hardware RAID5 does not meet the application performance requirements or mirrored data is an availability requirement, the next choice to be made is between hardware RAID1 or data host-based software RAID1. This choice depends not only on the availability requirements of the application, but on performance requirements as well.

Small sequential writes can incur performance degradation on the array partner group arrays due to cache mirroring. Thus, it may be better to use software mirroring for those applications. For other applications, the increased simplicity provided by hardware RAID1 is recommended. While there is no substitute for performance testing with actual application workloads, as a simple rule of thumb, use hardware RAID5 for applications with  $\geq 70\%$  READs or with large sequential writes. Use hardware RAID1 for applications with small random writes, and use host-based mirroring for small sequential writes, or for mission critical applications requiring true RAID 1+0 protection.

With any configuration that relies on multi-pathing software to achieve redundancy, AP or DMP must be used. AP can be used for networks, in addition to storage devices and supports Dynamic Reconfiguration (DR); however, the current version does not support FC-AL hubs (note that, for any single host configuration, the use of hubs is discouraged because they reduce performance and availability, while increasing complexity and MTTR). DMP supports hubs, but not network failover.

Since DMP has undergone significant field testing and has a large installed customer base, it will be used in the recommendations presented here. The same principles apply, regardless of the multi-pathing software and logical volume manager used.

## Configurations *Without* Software Mirroring (SW RAID1/1+0/0+1)

Configurations where host-based mirroring is not employed should use Sun StorEdge T3 Arrays configured as partner groups with RAID5 or RAID1 LUNs to eliminate SPOFs. Array LUNs should be configured with standby disks, unless full time data center personnel and on-site spare disks are available. Standby disks allow data to immediately be reconstructed to a spare drive following a disk failure. Disadvantages include added costs, loss of usable disk space, and the fact that two reconstruction penalties are incurred per disk failure (one to the standby drive and one from the standby to the replaced drive). Note that the second reconstruction from the standby disk to the new disk is a disk-to-disk copy, and is generally much faster. If personnel and on-site spare disks are always available, standbys should not be configured, and failed drives should be replaced as soon as detected.

LUNs should be configured one per array to reduce complexity. Recall the `blocksize` parameter should be set appropriately, based on application I/O characteristics, and is global for all LUNs within an array group. It cannot be changed without removing the existing LUNs. Before a LUN can be used by a host system, it must be created, initialized, and mounted. As an example, the following array commands result in the creation, initialization, and mounting of two RAID5 LUNs with standby disks on an array partner group:

```
T3/::<1>vol add ulv1 data uld1-8 raid 5 standby uld9
T3/::<2>vol init ulv1 data
T3/::<3>vol mount ulv1
T3/::<4>vol add u2v1 data u2d1-8 raid 5 standby u2d9
T3/::<5>vol init u2v1 data
T3/::<6>vol mount u2v1
```

The data initialization process takes approximately 20 minutes per LUN for volumes configured as stated on an array with 18 GB drives and a 64 K blocksize. The array LUN size, disk drive size, and blocksize all affect volume initialization times. It is extremely important to perform a data initialization for any RAID5 or RAID1 LUN prior to putting data on it, to ensure parity is correct and mirrors are in sync, respectively.

The array cache mode should be set to their defaults of `cache=auto`, `mirror=auto` to allow the arrays to switch the cache mode between `writebehind` and `writethrough`, based on array health. The array multi-pathing support should

be enabled by setting `mp_support=rw` to export multiple paths to the Solaris OE. Since DMP will be used, VxVM must be installed on the data host. The array LUNs should be configured under VxVM on the data host as simple stripes or concatenated volumes. Even though an array partner group is fully redundant, every array (and most every computer system) has a passive midplane that could theoretically fail. Also, an array partner group could experience a firmware error or multiple hardware failures that cause data inaccessibility. Therefore, data striping across multiple array partner groups should be minimized. Where possible, simple volumes built on single LUNs can be used to minimize recovery time, should this type of failure occur.

## Configurations *With* Software Mirroring (SW RAID1/1+0/0+1)

Configurations relying on data host-based software mirroring for protection should use Sun StorEdge T3 Array single brick arrays, since cache mirroring and multi-path support are not required to achieve redundancy. The array LUNs can be configured in hardware as RAID0, RAID1, or RAID5. If using RAID0, it may be preferable to use 2 LUNs per array—disks 1 through 4 in one LUN, and disks 5 through 9 in the other. This will minimize the amount of data resync required following an array disk failure, since any disk failure in a RAID0 LUN results in the loss of that LUN. As an example, the following array commands result in the creation, initialization, and mounting of two RAID0 LUNs on an array single brick array.

```
T3:/:<1>vol add ulv1 data uld1-4 raid 0
T3:/:<2>vol init ulv1 fast
T3:/:<3>vol mount ulv1
T3:/:<4>vol add ulv2 data uld5-9 raid 0
T3:/:<5>vol init ulv2 fast
T3:/:<6>vol mount ulv2
```

Note the initialization of RAID0 LUNs can be performed in `fast` mode, which takes approximately 5 seconds, versus the 20 minutes required for RAID5 or RAID1 LUNs, because there is no parity to generate or mirrors to synchronize. With RAID0, it is extremely important that disk replacement procedures (within the selected logical volume manager) are closely followed to prevent data corruption.

If using RAID5 or RAID1 array LUNs, they should be configured one LUN per tray, using the same guidelines for non-software mirrored configurations.

The array LUNs should be configured under VxVM on the data host as "Striped Pro" (RAID1+0) volumes with Dirty Region Logs (DRLs). This can take from minutes to hours, depending on the size of the volume and number of stripes.



Striped Pro (or RAID1+0) volumes maximize availability by mirroring at the LUN level instead of the plex level, a feature introduced with VxVM 3.0. They also reduce volume resync times by minimizing the data that is stale following a disk failure.

DRLs reduce volume resync time following a system crash, which can reduce restoration time to optimal system state. While DRLs contribute to improved availability, they can also reduce performance during normal operation if their locations in relation to data are not wisely chosen. Generally, separating DRLs on their own dedicated device yields the best performance, however other configuration and cost constraints may make this impractical.

All data mirrors should be on different arrays or partner groups, and should be housed in separate racks where possible. In this case, data stripes can span multiple arrays or partner groups.

The cache mode and mirror settings should be changed to `cache=writebehind`, `mirror=off`, since individual arrays are not relied on for redundancy. Even if the cache and mirror settings are left at their defaults of `auto`, the array firmware knows that in an array single brick configuration, cache mirroring is not possible. It allows the cache to stay in `writebehind` as long as the redundant FRUs (PCUs, disks, loop cards) are healthy. The `mp_support` variable should be left at its default setting (`none`), since multi-pathing with single arrays is not possible.

## Other Configurations

There are likely endless ways to combine software and hardware RAID on the Sun StorEdge T3 Array, but the two previous examples should cover most application requirements. One combination that is extremely resilient is hardware RAID5 array single brick LUNs with software RAID1+0. This configuration can achieve very high levels of availability, since it combines the protection of hardware RAID5 for individual disks with the availability of software mirroring for the controllers. It will likely figure prominently in multi-hosted or clustered configurations. Another valid (but not recommended) configuration employs the array single brick arrays with RAID5 or RAID1 LUNs without host-based mirroring. This configuration has SPOFs which may be acceptable in some workgroup environments, but should not be used in mission critical environments. Certainly, much more complex configurations can be developed. Remember, however, when designing storage solutions, simplicity is key.

---

# Monitoring

Sun StorEdge T3 Arrays, like any important system, require diligent monitoring to ensure failures are rapidly sensed, diagnosed, and corrected. The array has incorporated multiple monitoring and reporting facilities to aid the user in keeping their systems in an optimal state. Those that will be discussed include remote message logging, StorTools 3.3, Sun StorEdge Component Manager 2.2, SNMP management, and Sun™ Remote Services (SRS) software. Important tips for maintaining optimum PCU battery health are also included.

## Remote Logging

The Sun StorEdge T3 Array has a logging facility (`syslog` daemon) similar to the Solaris OE "`syslogd`" daemon. The logging facility can be configured to remotely log events to other hosts (log hosts) on the array network. The logging facility is controlled by the `logto` parameter, under the `set` command (`T3:::<1>set logto [1 | * | <filename>]`). A setting of `1` sends `syslog` messages to the console, a setting of `*` routes messages as specified in the array's `/etc/syslog.conf` file, and the `<filename>` setting causes messages to be logged locally on the array to the specified file. Normally, the setting should be `*`. The level of message reporting is controlled by another `set` parameter, `loglevel`. There are 5 `loglevel` settings, 0 through 4:

- 0 - No logging
- 1 - Errors
- 2 - Warnings and above
- 3 - Notices and above
- 4 - Informational and above

Normally, the `loglevel` should be set to 3 (`T3:::<1>set loglevel 3`).

To remotely log messages to a host system, the array `/etc/syslog.conf` and the `/etc/syslog.conf` file on the log host must be edited. Since the array operating system does not have a built-in editor, the file must be edited and then sent via FTP to the array. After the files have been edited, the array `syslog` facility must be restarted (`T3:::<1>set logto *`) and the log host `syslogd` must be restarted (`# kill -HUP <syslogd pid>`). The logging location specified in the log host system's `syslog.conf` file must be created prior to restarting `syslogd` if it does

not already exist (`# touch /var/adm/messages.t300`). Sample array `syslog.conf` and log host `syslog.conf` files are shown in FIGURE 3 and FIGURE 4, respectively.

```
# T3 syslog.conf
# facility.level action
# messages to local syslog file
*.info          /syslog
# messages to syslogd on another host
*.info          @<IP address>
```

**FIGURE 3** Sample Sun StorEdge T3 Array `/etc/syslog.conf`

```
# Log host syslog.conf
local7.info     /var/adm/messages.t300
```

**FIGURE 4** Sample Sun StorEdge T3 Array Log Host `/etc/syslog.conf`

Complete instructions for configuring the array logging facility are available in the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

As can be seen in FIGURE 3 and FIGURE 4, a setting of `info` was chosen for the facility reporting level, although the `loglevel` setting is 3—notice level. This is done so the files do not have to be edited to change the desired logging level. Instead, a simple change to `loglevel` on the array is all that is required to see more or less detail. All data hosts should be configured to log the attached array system messages to `/var/adm/messages.t300`. This will facilitate monitoring by StorTools. All management hosts should also be configured to log their managed arrays to `/var/adm/messages.t300`. For all log hosts, a method of routinely archiving the `/var/adm/messages.t300` file must be devised to prevent the file from growing without bound. A script to perform this function called `newt300log` is included with StorTools 3.3 patch 109090-04 and later. An example root crontab entry using this script is shown in FIGURE 5.

```
# T3 Log Archiver
10 3 * * 0 /opt/STORtools/bin/newt300log
```

**FIGURE 5** Sample Root crontab Entry

## StorTools 3.3

StorTools 3.3 software package provides monitoring, diagnostic, and installation support for Sun StorEdge T3 Arrays and other Sun StorEdge products. Only the monitoring features as they apply to the arrays are discussed here. StorTools software should be installed on all data hosts and kept up-to-date with required patches. See the *Sun StorEdge StorTools User's Guide* for complete installation instructions. StorTools software and manuals can be downloaded from the location given in the *References* section.

After StorTools software has been installed and the arrays configured to log to `/var/adm/messages.t300` on the data host, the monitoring features should be enabled. StorTools has three separate monitoring features, or checks. The first check consists of performing a scan of the current storage state, comparing it to a previously known state, and flagging any differences. The second check examines `/var/adm/messages` for storage related issues. The third check examines `/var/adm/messages.t300` for array errors and warnings. All checks work by sending email descriptions of problems to user configured email addresses.

To enable the checks, start StorTools software by executing:

```
# /opt/STORtools/bin/stormenu
```

as root on the data host. This performs a simple inquiry of the storage devices and reports PASS/FAIL. Press return to get to the main menu. Before enabling the checks, the storage subsystem's current state must be saved. This saved state is called a "Golden Snapshot." It is created by selecting the menu item "Create Configuration Snapshot File." Once the Golden Snapshot has been created, it should be updated following any changes to the storage hardware configuration or logical volume layout, to prevent the generation of spurious warnings.

To enable the first check, select "Storage Monitor" from the stormenu main menu, then select "Storage Status Check Setup," and finally, select "Enable Status Checking." The menu prompts will step through creating a crontab entry for root that will perform the check. It is recommended this check be performed three times daily. All checks should be configured to send email to the system administrator or other responsible party.

To enable the second check, select "Storage Monitor" from the stormenu main menu, then select "Storage Status Check Setup," and then "Enable StorEdge `/var/adm/messages` Checking." This check has thresholds for certain types of storage related error messages (e.g., ten loop "OFFLINE" messages per day on a given FC-AL loop) and works by scanning the past 24 hours of `/var/adm/messages`. It is recommended this check be performed hourly. If a problem is reported by this check and corrective action taken, it will continue to report the problem as long as the thresholds were exceeded in the past 24 hours. To prevent further reporting of the same problem, simply archive `/var/adm/messages`.

To enable the last check, select “Storage Monitor” from the stormenu main menu, then select “Storage Status Check Setup,” and then “Enable T300 syslogd /var/adm/messages.t300 Checking.” This check will capture and send alerts for all “Error” and “Warning” level messages forwarded to the monitoring host by the array syslog daemon. It is recommended this check be performed every 10 minutes with a 20 minute scan window. This means the /var/adm/messages.t300 file will be scanned every 10 minutes for events that occurred during the past 20 minutes in time. It is essential that the system time for the arrays and the data host system time be within 10 minutes of each other, or errors could be missed. Using a scan window twice the size of the scan frequency can cause the same error to be reported twice. However, it also should prevent errors from being missed due to small time differences between the arrays and the data hosts. An example crontab entry created by these menu selections is shown:

```
# StorTools Status Checking Entry
0 0,8,16 * * * /var/opt/STORtools/bin/storconfig -cmp -log -mail
root,sysadmin
# StorTools StorEdge Messages Checking Entry
0 * * * * /var/opt/STORtools/bin/mess_alert -f /var/adm/messages
-mail root,sysadmin
# StorTools syslogd Checking Entry
0,10,20,30,40,50 * * * * /var/opt/STORtools/bin/mess_alert -f /
var/adm/messages.t300 -mail root,sysadmin -h 0 -m 20
```

## Sun StorEdge Component Manager 2.2 Software

Sun StorEdge Component Manager 2.2 software provides monitoring and management for Sun StorEdge A5x00 and Sun StorEdge T3 Arrays. It provides a graphical user interface that displays the status of the arrays. The features applicable to the arrays include alarm notification, remote reporting, viewing of FRU status and properties, physical subsystem views, LUN configuration, and offline diagnostics.

As stated previously, Sun StorEdge Component Manager software should be installed on the management host. To enable monitoring and management of the arrays, simply add their IP addresses to the /etc/opt/SUNWesm/mo/hosts file and restart the appropriate daemons by executing /usr/opt/SUNWesm/sbin/esm\_orderly stop followed by /usr/opt/SUNWesm/sbin/esm\_orderly start. Sun StorEdge Component Manager software works by querying HTML files located in each array’s local /web directory. These files must be kept up-to-date, along with the array firmware, for Sun StorEdge Component Manager to work properly. Sun StorEdge Component Manager should be configured to log errors and send email alerts to system administrators as part of a comprehensive monitoring plan. To facilitate viewing of

the array messages produced during diagnostics, the array can be configured to push the messages via HTTP by appending the following to the array `syslog.conf` file:

```
#CM 2.2 Diagnostic Message Push
*.info | http_push
```

Complete instructions can be found in the *Sun StorEdge Component Manager 2.2 User's Guide*.

## SNMP Management

The Sun StorEdge T3 Array can provide remote notification of array events to designated hosts using Simple Network Management Protocol (SNMP) via a Management Information Base (MIB) stored on its local file system. The array MIB is located at `/web/snmp/t300.mib` on the array local file system. Like the HTML files used by Sun StorEdge Component Manager software, the MIB must be kept up-to-date to ensure accurate reporting to the SNMP management software. The latest MIB is available in patch 109115, along with all other files and firmware needed for the array. To send SNMP traps to a management host, edit the array `/etc/syslog.conf` file appropriately. As an example, adding the following to the array's `/etc/syslog.conf` file will send information messages and higher in the form of SNMP traps to the desired host (depending on the setting of the `loglevel` parameter):

```
#SNMP Trap Configuration
*.info | snmp_trap <IP address>
```

Remember to restart the `syslog` facility by executing `T3:/:<1>set logto *` for the changes to take effect. For more information, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*, and the chosen SNMP management software documentation.

## Sun Remote Services

Sun Remote Services (SRS) is a suite of services provided by Sun to help customers monitor and manage the systems vital to their business. SRS helps improve system availability through proactive problem detection. They also bring advanced, full time remote monitoring and management capabilities to customer's IT environments, and can provide fast, direct access to Sun's engineers and IT experts.

These services are especially useful for enterprises deploying multiple Sun StorEdge T3 Arrays in a mission critical environment. The array is SRS ready, meaning there are SRS modules that have been written to scan for specific errors from an array and prompt the SRS monitoring personnel to take corrective action. More information on deploying SRS can be obtained from your local Sun service representative.

## Power/Cooling Unit Batteries

Each PCU contains a NiMH battery that flushes any uncommitted write cache data to disk in the event of power loss. It is important to maintain the battery properly to ensure this capability remains intact. The batteries have an installed life span of two years, after which they must be replaced by replacing the PCU. The Sun StorEdge T3 Array will regularly log reminder messages about battery replacement, starting 45 days before the expiration date. After the battery life expires, the cache is placed in `writethrough` mode. In order to maintain battery health during its normal life span, it must be cycled by discharging and recharging at regular intervals. This function is accomplished by the battery scheduler daemon, which runs as part of the array operating system. This daemon is controlled using the `refresh` command:

```
T3:/:<8>refresh
usage:
  refresh -c --> start refresh
  refresh -s --> get status
  refresh -i --> re-init command file
  refresh -k --> kill current refreshing task
```

A battery refresh is performed in two phases, a discharge phase and a recharge phase. Each PCU must complete these two phases serially to ensure write cache data is protected during the refresh cycle. In configurations with multiple arrays cabled together, PCU1 on each array performs the phases in parallel followed by PCU2. During the discharge phase, the cache is placed in `writethrough`. This phase typically lasts 6 to 20 minutes, and is recorded in the `syslog` as the battery hold time. The minimum hold time for a healthy battery is 6 minutes.

Following this, the recharge phase lasts up to 12 hours. Thus, during a normal refresh cycle, the cache will revert to `writethrough` twice, once for PCU1 and once for PCU2. Refreshes are required to be performed every 28 days. This means that over a four-week period, the cache will be in `writethrough` for up to 40 minutes for battery refreshes. While this typically represents an insignificant performance penalty, the time during which it occurs should be scheduled to prevent it from happening during peak usage periods, or from being misdiagnosed as a FRU problem.

The battery refresh daemon has a configuration file, `/etc/schd.conf`, that can be used to specify the start time and refresh period. Simply edit the file, FTP it to the array, and execute `refresh -i` to reinitialize the daemon. The daemon works on GMT, so the offset between the current time zone and GMT must be added when specifying the refresh start time in the file. An example `schd.conf` file is shown in Figure 6.

```
BAT_BEG MM-DD-YYYY, hh-mm-ss
BAT_CYC 28
```

FIGURE 6 Sample Sun StorEdge T3 Array `/etc/schd.conf`

---

## Summary

### Guidelines for Maximizing RAS for all Sun StorEdge T3 Array Configurations:

- Follow the rack and table top mounting requirements to ensure adequate cooling is maintained.
- Power the arrays from separate circuits.
- Maintain proper battery hold times and refresh cycles.
- Don't remove failed array FRUs unless a replacement is immediately available.
- Setup remote logging of the array messages to both application and management hosts.
- Setup StorTools software monitoring features on the application hosts.
- Setup Sun StorEdge Component Manager software to monitor the arrays on the management host.
- Setup a TFTP host and maintain the appropriate T3 firmware files.
- Contract with Sun Enterprise Services to deploy SRS.
- Minimize use of FC-AL hubs where possible.
- Keep array firmware and related data host patches up-to-date within the confines of a well-defined change control process.
- Connect redundant paths, whether used for multi-pathing or host mirroring, to separate host adapters, system buses, I/O boards, and system boards where applicable.
- Configure active paths over separate system buses to maximize bandwidth.



## Guidelines for Maximizing RAS for Sun StorEdge T3 Array Configurations not Employing Software Mirroring:

- Use the arrays configured as partner groups.
- Configure the array LUNs as RAID1 or RAID5 with 1 LUN per array.
- Set the array cache mode to `cache=auto`, `mirror=auto`, or do not use writebehind caching.
- Set the array `mp_support=rw` to export multiple paths to Solaris OE.
- Configure the array standby disks, unless full time data center personnel are available.
- Use VxVM and DMP. Create simple striped or concatenated volumes.
- Minimize striping of data host volumes across multiple partner groups.

## Guidelines for Maximizing RAS for Sun StorEdge T3 Array Configurations Employing Software Mirroring:

- Use arrays configured as single bricks.
- Configure T3 LUNs as RAID5, RAID1, or RAID0.
- Configure RAID5 or RAID1 LUNs with 1 per array.
- Configure RAID0 LUNs with 2 per array, 4 disks in one LUN, and 5 disks in the other.
- Set the array cache mode and mirroring to `cache=writebehind`, `mirror=off` to maximize performance.
- Leave the array variable `mp_support=none`, since multi-pathing is not available in single brick configurations.
- Use VxVM and DMP. Create RAID1+0 data volumes with DRLs.
- Place data mirrors on separate bricks, in separate cabinets where possible.

## Guidelines for Maximizing RAS for Sun StorEdge T3 Array Single Brick Configurations Employing Multiple Arrays Without Host-Based Software Mirroring (*Not a Recommended Configuration*):

- Configure the array LUNs as RAID5 or RAID1 with 1 LUN per array.
- Configure the array standby disks, unless full time data center personnel are available.
- Do not span data host volumes across multiple arrays.

- Set the array cache mode to `cache=writethrough, mirror=off`, or assume the additional risk associated with non-mirrored writebehind caching.
- Do not use in any mission critical environment.

---

## References

The following recommended reading is available from the web pages listed below:

[http://www.sun.com/products-n-solutions/hardware/docs/Network\\_Storage\\_Solutions/Midrange/T3ES/index.html](http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Midrange/T3ES/index.html)

*Sun StorEdge T3 Disk Tray Site Preparation and Planning Guide*

*Sun StorEdge T3 Disk Tray Installation, Operation and Service Manual*

*Sun StorEdge T3 Disk Tray Administrator's Guide*

*Sun StorEdge T3 Disk Tray Installation Task Map*

*Sun StorEdge T3 Disk Tray Cabinet Installation Guide*

*Sun StorEdge T3 Configuration Guide*

*Sun StorEdge T3 Disk Tray Release Notes*

[http://www.sun.com/products-n-solutions/hardware/docs/Network\\_Storage\\_Solutions/Cabinets\\_and\\_Enclosures/index.html](http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Cabinets_and_Enclosures/index.html)

*Sun StorEdge Expansion Cabinet Installation and Service Manual*

*Sun StorEdge Expansion Cabinet Fan Tray Installation*

*Sun StorEdge Expansion Cabinet Front Screen Door Installation Guide*

*Rackmount Placement Matrix*

[http://www.sun.com/products-n-solutions/hardware/docs/Software/Storage\\_Software/Sun\\_StorEdge\\_Component\\_Manager/index.html](http://www.sun.com/products-n-solutions/hardware/docs/Software/Storage_Software/Sun_StorEdge_Component_Manager/index.html)

*Sun StorEdge Component Manager 2.2 Installation Guide*

*Sun StorEdge Component Manager 2.2 User's Guide*

*Sun StorEdge Component Manager 2.2 Release Notes*

<http://www.sun.com/microelectronics/vts/download/sunvts.html>

*SunVTS User's Guide*

<http://reseller.sun.com/tools/stortools/>

*Sun StorEdge StorTools User's Guide Version 3.3*

*Sun StorEdge StorTools Release Notes Version 3.3*

<http://www.sun.com/storage/wp.html>

*Sun StorEdge T3 Array: A Technical White Paper*

---

*[Author's Bio: Ted Gregg](#)*

*Ted Gregg has been with Sun for 3 years. He is currently an engineering manager on the Network Storage team within Customer Problem Resolution Engineering. He works primarily as a liaison between development engineering and the field organizations, providing training, improving troubleshooting methods, determining best practices, and working on critical customer problems.*