

Tech Company

Zen4 Software Configuration Guide

Tech Company Systems, Inc. Legal Notice

Copyright © 2005, 2006 Tech Company Systems, Inc. All rights reserved.

Tech Company Systems, Inc.

11 Lane, Suite 222

Huntsville, AL 11222

Phone: 555-12-1212

www.TechCompany.com

Disclaimer

The information contained in this publication is subject to change without notice. Tech Company Systems, Inc. makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Tech Company Systems, Inc. shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

Trademarks

The Tech Company logo, Tech Company, Z-Series, Z-Visor, and all other Tech Company Systems, Inc. product names and service marks are trademarks of Tech Company Systems, Inc.

AMD Opteron is a registered trademark of Advanced Micro Devices. Linux is a registered trademark of Linus Torvalds. RED HAT is a registered trademark of Red Hat, Inc. SUSE is a trademark of Novell. VMware® is a registered trademark of VMware, Inc. Microsoft and Windows are registered trademarks of Microsoft Corporation. IBM and Tivoli are registered trademarks of IBM Corporation.

Other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Printed in USA.

Contents

Chapter 1 Zen4 System and Hardware Overview	4
The following sections are included in Chapter 1:	4
System and Architectural Overview	5
The following sections are included in the System and Architectural Overview:	5

Fabric Computing	5
Functional Features	5
Functional Overview	6
Integrated Computing and Networking	6
Resource Sharing Across a Switch Fabric	6
Switch	7
Network Interfaces	7
Virtual I/O Interfaces	8
Flexible Provisioning	8
Technical Overview	9
The following sections are included in the Technical Over- view:	10
Processors Memory and Partitioning	10
System Management	11
System I/O Fabric	11
Boot Support	12
Backplane	13
Internal Hard Drive	13
System Battery and Time	13
System Battery	13
System Time	13
Software Components	13
Typical System Installation	15
Zen4 Hardware Overview	16
The following sections are included in the Hardware Over- view:	16
Processor and Memory Modules	17
Z-Series System Module (ZSM)	17
Ethernet Network Module	17
Fibre Channel Modules	18
Application Service Modules	18
Hot-Plug and Controlled Hot Removal Capability	18
High Availability	18
Virtual I/O Controller (VIOC) and Virtualization Cap- abilities	19

The following sections are included in Virtual I/O Controller (VIO) and Virtualization Capabilities:	19
Virtual Network Interface Controllers	19
Virtual Storage Interfaces	20
Partitioning and Virtualizing Resources	21
Data Flow in the Zen4 System	22
Appendix A: Product Support	23
Warranty	24
Verifying Warranty Status	25
Before you begin	25
To verify your warranty status	25
Result	25
Service Features	26
Service Levels	27
Priority Levels	27
Contacting Customer Support	28
Replacing a Component	29
Before You Begin	29
Replacing a Component	29
Result	29
Returning a Component	30
Before You Begin	30
Returning a Defective Component	30
Result	30
Ordering Replacement Parts or Spares	31
Before you begin	31
Ordering replacements parts or spares	31
Result	31
Glossary	32
Index	33

Chapter 1 Zen4 System and Hardware Overview

Chapter 1: Zen4 System and Hardware Overview introduces the Tech Company Zen4 server and describes its relationship to advanced data center computing. It also gives an overview of important network concepts in the Zen4 system and provides a brief overview of the Zen4 hardware components.

The following sections are included in Chapter 1:

- [System and Architectural Overview](#)
- [Technical Overview](#)
- [Typical System Installation](#)
- [Zen4 Hardware Overview](#)
- [Virtual I/O Controller \(VIOC\) and Virtualization Capabilities](#)
- [Partitioning and Virtualizing Resources](#)
- [Data Flow in the Zen4 System](#)

System and Architectural Overview

The Tech Company Zen4 system integrates highly redundant multiprocessing servers and networking technology in a modular design, which provides efficient use of space and simplifies management tasks. The Zen4 system provides improved data center performance by integrating resources and sharing key components to provide high density processing and Layer 2 network connectivity.

The Zen4 system is architected from the ground up to exploit a tight coupling between the network data plane and the application computing plane.

The following sections are included in the System and Architectural Overview:

- [Fabric Computing](#)
- [Functional Features](#)
- [Functional Overview](#)
- [Integrated Computing and Networking](#)
- [Resource Sharing Across a Switch Fabric](#)
- [Switch](#)
- [Network Interfaces](#)
- [Virtual I/O Interfaces](#)
- [Flexible Provisioning](#)

Fabric Computing

Fabric computing is the logical extension of the existing concept of a VLAN to include servers and services. Zen4 system enables the convergence of the network and application networking environment. Fabric computing is the mechanism for achieving this integration.

Tech Company Systems offers the system administrator the opportunity to unify network and server management. Typically four different devices are used to implement the network, servers, storage, and application services. The Zen4 system unifies these four components into one server. The Tech Company Zen4 system management infrastructure manages VLANs as a transparent extension of a management fabric.

Functional Features

The Zen4 system, built around an innovative terabit-class fabric, uses proprietary virtualization techniques. The servers combine partitionable SMP complexes with high-performance networking and storage interfaces to deliver world-class performance. Volume components (64 bit x86

architecture) and off-the-shelf operating systems provide binary compatibility for existing applications.

Zen4 hardware resources combine to provide an enterprise server with the following features:

- **Performance** for network-intensive applications through the tight coupling of network line cards and processor complexes
- **Scalability** through the modular multi-chassis architecture
- **Manageability** through the consolidation of server and network resources, simplified provisioning and management, virtualized server capabilities, and location-independent server deployment
- **Resiliency** for enterprise applications through server partitioning and integrated load-balancing
- **Interoperability** with existing application infrastructure through support of standards like TCP/IP, 1-Gbps/10-Gbps Ethernet, 2-Gbps Fibre Channel, and SNMP

Functional Overview

The Zen4 system is a fabric server that integrates the following data processing and I/O capabilities:

- Partitionable N-socket symmetric multiprocessing (SMP) with large memory using industry-standard x86-64 microprocessors and Linux® operating systems
- Over 30-Gbps Ethernet network and Fibre Channel I/O bandwidth
- Application services such as server load balancing, and Secure Socket Layer (SSL)

Integrated Computing and Networking

The Zen4 system can be deployed to deliver enterprise applications and services on demand from resources within a single chassis, across multiple chassis in one location, or over geographically dispersed sites. The Zen4 system is also integrated with an industry standard Ethernet switch fabric to seamlessly integrate state-of-the-art computing resources with high-speed networking and I/O.

Resource Sharing Across a Switch Fabric

The Zen4 system combines eight-socket partitionable SMP using 64-bit AMD Opteron® processors with 1-Gbps/10-Gbps Ethernet networking, 2-Gbps Fibre Channel storage I/O, and offloaded application or acceleration services. Hardware-based partitioning and virtualization enable pooling and sharing of resources, as seen in [Figure 1-1. Multi-chassis Fabric with Resource Sharing](#).

Figure 1-1. Multi-chassis Fabric with Resource Sharing

The network and storage I/O can be dynamically allocated across a built-in 80-Gbps switch fabric. Fabric-wide Quality of Service (QoS) ensures application service levels in a shared environment. Multiple Zen4 chassis can be directly coupled and managed as a single pool of resources.

The highly integrated fabric server collapses multiple computing and network infrastructure tiers and enables fabric computing across multiple sites. Policy based, fabric-wide management software allows rapid deployment, provisioning, and scaling of resources from any point on the fabric. The result is a mainframe-class architecture that offers high throughput and high stability for Linux and Windows applications.

Switch

The Zen4 system includes a built-in switch. This switch behaves and can be configured just like any state-of-the-art Layer 2 switch. It supports virtual LANs (VLANs) and the set of Layer 2 protocols.

The switch supports physical and logical ports. Logical ports are used to create Layer 2 connections to virtual network interface controllers (VNICs) within the system. Logical ports can also be used to create Layer 3 ports on a VLAN to allow IP routing to and from the VLAN. Logical ports are managed in exactly the same way that physical ports are managed.

Any port on the switch, logical or physical, can be configured to behave as a Layer 2 port. As a result of the IP address assignment, an underlying MAC address will also be assigned, allowing this IP assignment to be addressed and packets to be sent to it. This IP address is generally called a “gateway” IP address.

Network Interfaces

Network interfaces provide the means for external networks to be connected to the node by providing the interface between the external network and the internal IO fabric. A network interface provides all of the functionality of a typical enterprise switch, including classification, queuing, routing, policing, traffic shaping, etc. Network interfaces are provided currently for Ethernet (1-Gbs and 10-Gbs) and 2-Gbps Fibre Channel.

Ethernet network interfaces are deployed on modules called Ethernet Network Modules (ENMs). The ENMs contain the MACs and PHYs for the particular media, NPUs, traffic managers, support processors, and a fabric interface.

Fibre channel interfaces are deployed on modules called Fibre Channel Modules (FCMs). FCMs contain the MACs and PHYs for the Fibre Channel media, support processors, and a fabric interface.

Virtual I/O Interfaces

Tech Company virtual I/O interfaces conform to the following principles:

- A fabric server is built from a partition consisting of a processor and memory complex to which virtual I/O interfaces are added.
- There are three types of virtual I/O interfaces:
 - VNICs, which present Ethernet interfaces to the servers within the context of VLANs.
 - iSCSI initiators, which offload the iSCSI protocol to hardware within the context of storage VLANs.
 - Virtual service interfaces (VSIs), which provide data network connectivity to Application Services Modules (ASMs)
- All virtual I/O interfaces offer independent bandwidth control and three available classes of service (managed as priorities): low, medium, and high.

Flexible Provisioning

The Tech Company Zen4 unified fabric management architecture merges the network control plane with server provisioning and management, bringing simplicity, agility, and cost-efficiency to the operation of the servers.

Technical Overview

The Zen4 system is a collection of resources arranged around a coherent fabric and a network or I/O fabric. Three broad types of resources are attached to the fabrics and linked through an internal switch: processing and memory, I/O modules (Ethernet and Fibre Channel), and application off-load or acceleration services. The Z-Series System Module (ZSM) monitors and manages these resources. The processors and associated memory are organized as standard SMP servers, allowing flexible combining and partitioning to meet application needs. The Ethernet and Fibre Channel interfaces provide connections to the external network and to storage resources. The Zen4 system includes built-in Layer 2 network switching capabilities.

Within a Zen4 system are multiple operating environments, including user operating environments for customer application deployment and internal operating environments for the Zen4 software. The user operating system (for example, SUSE® Linux Enterprise Server 9 SP2) runs in a partition to create a server that runs user applications.

Reliability, availability, and serviceability (RAS) are three qualities that represent the collective quality-of-service attributes of network systems. The Zen4 system has a combination of features, hardware design, and software design that maximizes uptime and minimizes downtime. Hardware design features that contribute to high levels of RAS include front access to most components for servicing, redundant cooling and power systems, and redundancy of system control modules and network interfaces. A graphical user interface (GUI) and command line interface (CLI) are available for managing system resources; in addition, an SNMP agent allows the Zen4 chassis to be integrated into enterprise monitoring systems such as HP OpenView or IBM Tivoli® for diagnostics, monitoring, and troubleshooting. The partition boot architecture of the Zen4 system helps minimize downtime arising from disruptive operating system updates.

Server administration and network administration are combined in the Zen4 system; the flexible security model provided by the Zen4 platform allows these functions to be separated in settings where this is preferred. The Zen4 platform provides fine-grained security and role-based access control that defines specific access privileges for different user roles. In addition to the default user roles, you can also define other specific roles and access privileges.

[Figure 1-2. Zen4 System Block Diagram](#) shows a block diagram of the Zen4 system.

Figure 1-2. Zen4 System Block Diagram

The following sections are included in the Technical Overview:

- [Processors Memory and Partitioning](#)
- [System Management](#)
- [System I/O Fabric](#)
- [Boot Support](#)
- [Backplane](#)
- [Internal Hard Drive](#)
- [System Battery and Time](#)
- [Software Components](#)

Processors Memory and Partitioning

Application processors are deployed on modules called Processor and Memory Modules (PMMs). PMMs support general-purpose computing requirements and are available in two types: PMM-04xx series and PMM-02xx series. Currently, the PPMM-0410 and PMM-0200 are supported.

- PMM-0410s provide CPU sockets for up to four AMD Opteron 64-bit x86 CPUs with memory. The PMM-0410 also offers a total of 8 Gbps of storage I/O bandwidth using Fibre Channel interfaces. The Zen4 chassis supports two installed PMM-0410s.
- PMM-0200s provide CPU sockets for a maximum of two AMD Opteron 64-bit x86 CPUs with memory. Applications can be installed on a PMM-0200 to provide low-latency access by the PMM-0410. Such a configuration can improve performance for certain applications on the PMM-0410 that depend on other applications that can be accessed with low latency.

The application processors and memory are the only resources that attach to the coherency fabric. The processors and memory on a PMM can be organized into standard SMP servers; in a chassis equipped with PMM-0410s, multiple PMMs can be linked through the coherency fabric (backplane) to create instances of SMP servers containing up to eight processor sockets. Smaller partitions of two or four processors can also be configured.

A chassis equipped with two PMM-0410s can support eight system processors (or 16 system processors with dual-core processing), 128 GB of system memory, 30 Gbps allocated between network traffic and storage traffic, and three Application Service Modules (ASMs).

The PMM-0410s and the coherency fabric support physical partitioning of the processors and memory. Because each partition is an independent server, multiple independent instances of operating systems can be hosted concurrently. Both types of partitions allow a partition to be created, initialized, started, stopped, reset, or destroyed without affecting

the operation of other partitions. Each partition appears as a separate entity for system management.

Physical partitioning is implemented in the hardware and in system software. The hardware ensures that hardware and software faults in one hardware partition cannot propagate to another partition.

For more information about the PMMs, refer to the Zen4 Hardware Guide (PIN 502-0001-01).

System Management

System management is implemented through a combination of management software components and a distributed hardware management subsystem. Through the hardware management subsystem, system components communicate with the ZSM, which manages chassis resources. The hardware management subsystem also allows remote, "lights-out" operation of the Zen4 system. This management capability extends to a system in the standby power state, when it has limited operational capabilities as it waits for instructions. All system messages are directed to the ZSM console and system log (syslog) facility.

Access to the Zen4 system for management and software installation is through network or serial connections, eliminating the need for keyboard, video, and monitor connectors on the system itself. In addition, because all necessary software installation can be accomplished through a network connection, the Zen4 system is designed without a CD-ROM drive.

The Zen4 chassis has an LCD front panel display with associated selection buttons. Using the front panel display, you can view system status information and restart the chassis. The CLI, accessible through a direct console port connection or a network connection, provides complete management functions for the system. The Tech Company Z-Visor™ software provides the same management capabilities as the CLI, through a user-friendly graphical interface. Z-Visor also provides visual alarms and alerts.

The management network is backed up by an intra-chassis management bus based on RS-485 serial links.

System I/O Fabric

All system modules are tied together through an I/O fabric implemented as a 32 x 32 crossbar switch. The system has 11 network fabric ports, allocated as follows:

- Two 10-Gbps ports are assigned to each of two PMM-0410s (up to four ports total).
- One 10-Gbps port is assigned to each I/O module installed in slot 5, 6, or 7 of the Zen4 chassis (support for up to three ENMs or up to two FCMs, providing up to three or two ports total).

- One 2.5-Gbps port is assigned to each ASM or PMM-0200 installed in slots 8, 9, and 10 of the Zen4 chassis (up to three ports total).
- One 2.5-Gbps port is assigned to the ZSM.

The switch hardware is instantiated on the ZSM.

ENMs are available with fixed 1-Gbps copper Ethernet ports, 1-Gbps SFP Ethernet ports, or 10-Gbps optical Ethernet ports. The ENMs provide an interface between the external network and the internal I/O fabric. Up to three ENMs can be installed in the Zen4 chassis, for a total network bandwidth of 30 Gbps.

FCMs are available with eight 2-Gbps Fibre Channel interfaces. One FCM provides access to an aggregate of 16-Gbps Fibre Channel I/O bandwidth through 2-Gbps Fibre Channel interfaces. A system with two installed FCMs, therefore, provides a total of 32-Gbps of Fibre Channel I/O. In addition, each PMM-0410 has dedicated Fibre Channel ports that are associated only with the server partitions on that PMM.

The fabric supports four priorities for general traffic. The use of the priorities is determined by software through configuration of the hardware, with the exception that all I/O control protocol read response messages are fixed to use fabric priority 0.

Boot Support

Traditional servers include local (built-in) or direct-attach disks to store an operating system that they boot from. This model imposes severe operational constraints that reduce flexibility in usage and prevent easy redeployment of computing resources.

The Zen4 system model is based on an abstract server definition that is not locked into a chassis. A server can be instantiated using resources available anywhere on the fabric and with resources contributed by multiple chassis. The operating system itself is a resource on the fabric, located on a network element or on a Fibre Channel storage array.

In this distributed model of a server, booting can occur in the following ways:

- Using the PXE mechanism, from a boot server in the netv.ork
- Using Fibre Channel:
 - From a storage area network (SAN) with the operating system installed on a logical unit number (LUN) on a Fibre Channel storage array that is connected to a SAN switch through a PMM.
 - From a direct-attached Fibre Channel storage array with the operating system installed on a LUN in the array that is attached directly to a Fibre Channel port on a Zen4 system
- Using iSCSI from an external storage target connected to an FCM in a Zen4 system

- Using iSCSI from an external iSCSI storage target connected to an ENM in a Zen4 system

Backplane

The Zen4 backplane is designed to allow system growth:

- Each slot is provided with connectors to support twice the bandwidth necessary for the initial system.
- By changing the ZSM, a system can be upgraded to twice the traffic on the backplane.
- The backplane carries two major fabrics, the system IO fabric, and the processor coherency fabric.

Internal Hard Drive

The Zen4 system is designed to provide "discless operation" and therefore does not include an internal hard drive for the user operating system (UOS). The Zen4 provisioning facility allows users to create flexible server configurations that can be revised and reconfigured quickly based on a number of considerations. This flexibility is lost if the operating system resides on an internal hard drive.

Long-term storage for the management application is provided by an on-board flash disk and a 2.5-inch hard disk on the ZSM. The flash disk is used for Zen4 operating system boot, configuration data, and application software. The 2.5-inch hard disk is used for logging.

System Battery and Time

System Battery

Over time, batteries can leak and damage other system components. To increase system reliability, the ZSM uses a capacitor, rather than a battery, to maintain the system time when the chassis is not plugged into a power source.

System Time

The system time can be maintained for approximately 48 hours while the chassis is unplugged.

Software Components

Software components of the Zen4 system include the system software image, UOS, and management software, which includes the CLI and the Z-Visor management interface.

The system software image resides on the ZSM; it manages the chassis hardware resources, including the server partitions. User operating systems are associated with partitions and execute on the processors on

the PMM. The CLI is accessible through a console connection or from a client workstation. Z-Visor software is accessed from a management workstation and manages the Zen4 sytem through an Ethemet connection to the ZSM.

Typical System Installation

The typical installation of the Zen4 system with four server partitions can be seen in [Figure 1-3. Logical View of Typical Installation](#). Each partition has two Fibre Channel I/O connections to a storage area network (SAN) and two virtualized Ethernet connections to the external network. The ZSM has an Ethernet connection and a serial connection to a management network.

Figure 1-3. Logical View of Typical Installation

Zen4 Hardware Overview

The Zen4 system consists of a 21U rack-mount chassis with the following installed resources:

- **Processor and Memory Modules (PMMs)** support general-purpose computing requirements and come in two series: PMM-0410s and PMM-0200s.
 - **PMM-0410:** PMM-0410s provide CPU sockets for up to four AMD Opteron 64-bit x86 CPUs with memory. The PMM-0410 also offers a total of 8 Gbps of storage I/O bandwidth using hard-wired Fibre Channel interfaces. PMM-0410s are installed in Zen4 chassis slots 1 and 2.
 - **PMM-0200:** PMM-0200s provide CPU sockets for a maximum of two AMD Opteron 64-bit x86 CPUs with memory. The PMM-0200 augments the CPU and memory resources already available through the two PMM-0410s. Each PMM-0200 represents an extra, isolated, two-socket server partition that can host certain proximal distributed application components that might be effectively off-loaded from the PMM-0410s. PMM-0200s are installed in chassis slots 8, 9, or 10.
- **Q-Series System Modules (ZSMs)** that implement provisioning, system management, and fabric management capabilities. The ZSM is installed in chassis slot 4.
- **Ethernet Network Modules (ENMs)**, each providing an aggregate of 10-Gbps of network I/O bandwidth, using either copper or fiber optic interfaces. Up to three ENMs can be installed in chassis slots 5, 6, and 7. (These slots can be used for either ENMs or FCMs.)
- **Fibre Channel Modules (FCMs)** that interface to storage arrays and storage devices, either directly or indirectly through storage area networks (SANs). An FCM can be installed in chassis slot 5, 6, or 7. (These slots can be used for either ENMs or FCMs.)
- **Application Service Modules (ASMs)** that provide acceleration facilities for specialized functions such as load balancing, encryption, and XML processing. Up to three ASMs can be installed in chassis slots 8, 9, and 10.

Multiple chassis can be interconnected to operate in a network fabric where physical resources in one chassis are available as virtualized resources and services anywhere on the fabric.

The following sections are included in the Hardware Overview:

- [Processor and Memory Modules](#)
- [Z-Series System Module \(ZSM\)](#)
- [Ethernet Network Module](#)
- [Fibre Channel Modules](#)

- [Application Service Modules](#)
- [Hot-Plug and Controlled Hot Removal Capability](#)
- [High Availability](#)

Processor and Memory Modules

Zen4 PMMs use AMD Opteron 800 Series processors. Four processors are available on each PMM-0410; two processors are available on each PMM-0200. HyperTransport™ links between the processors in a PMM and between two PMM-0410s allow all eight processors to be configured into a single symmetric multiprocessing (SMP) server. Smaller partitions of two or four processors can also be configured.

Z-Series System Module (ZSM)

The Q-Series System Module (ZSM) is the heart of the Zen4 system. The ZSM manages the resources, monitors the health of all the components, and takes appropriate actions, in case of any failures, to maintain the system stability. The module is designed to be replicated in the Zen4 system so that a failure on one module can instigate a failover to the other module. In general, a failure in any one subsystem on the module that causes that subsystem to failover will cause all subsystems to failover. This allows the failed module to be removed from the system for replacement.

Long-term storage for the management application is provided by an on-board flash disk and a 2.5-inch hard disk. The flash disk is used for OS boot, configuration data, and application software. The 2.5-inch hard disk is used for event logging. Software is responsible for maintaining a valid replica of critical data structures on the flash disk on the standby ZSM so that it is prepared to take over in the event of a failure. The storage devices on the ZSM are not customer-replaceable units.

The system clock on the ZSM provides system time-of-day.

Ethernet Network Module

The Ethernet Network Module (ENM) provides an interface between the external network and the internal I/O fabric. ENMs have either optical or copper ports that provide up to 10-Gbps of aggregate network I/O bandwidth. Up to three ENMs can be installed in the Zen4 chassis, for a total network bandwidth of 30-Gbps.

Three versions of the ENM are available:

- The model ENM-1011 has ten 1-Gigabit copper (RJ-45) ports.
- The model ENM-1012 has ten 1-Gigabit small form-factor pluggable ports (SFPs).
- The model ENM-1103 has one 10-Gigabit optical port implemented as an installed XENPAK optical transceiver.

- The model ENM-2411 has 24 1-Gigabit copper (RJ-45) ports.

Any ENM type can be installed in one of the Zen4 chassis slots 5, 6, or 7. The Zen4 system detects the installed module.

Fibre Channel Modules

Fibre Channel Modules (FCMs) interface to storage arrays and other storage devices, either directly or indirectly through Fibre Channel switches in storage area networks (SANs). Up to two FCMs can be installed in chassis slots 5, 6, or 7.

Application Service Modules

Application Service Modules (ASMs) provide hardware acceleration facilities for specialized functions such as load balancing, Secure Socket Layer (SSL), and XML processing. ASMs are installed in Zen4 chassis slots 8, 9, or 10.

Hot-Plug and Controlled Hot Removal Capability

All the modules used in the Zen4 system have circuitry to allow removal and insertion of a module without interrupting system power. This is accomplished by controlling the inrush current and power sequencing on the boards.

High Availability

The Zen4 system is designed with redundant subsystems to ensure high availability of resources. Four power modules that can be connected to two independent power feeds share the power load during normal operation. If one power feed fails, the remaining two power modules can provide sufficient power for the system. Two fan trays, each with six fans, provide cooling for the chassis. If one tray (or part of a tray) fails, the ZSM increases the speed of the remaining fans to maintain cooling until the faulty fan tray can be replaced. Multiple interfaces are available on the PMM-0410s and ENMs to provide redundancy for storage and network access.

Virtual I/O Controller (VIOC) and Virtualization Capabilities

Each PMM is assigned a virtual I/O interface with appropriate properties that depend on whether it is used for network, storage or services. A virtual I/O interface is called a virtual NIC (VNIC) when it is used for network data, or a VNIC used as iSCSI initiator when it is used for storage data. A virtual I/O storage interface is called a virtual service interface (VSI) when it is used in an ASM.

Each PMM-0410 has two VIOCs, and each PMM-0200 has one VIOC. In addition, each ASM and FCM has one VIOC. The VIOC implements the virtual NICs (VNICs) and iSCSI initiators that allow a server to communicate across the fabric with other servers within a complex, or with any server outside a complex. Through the VIOC, the management subsystem creates and manipulates these virtual interfaces. Virtual I/O interfaces support three types of traffic: network, storage, and application services.

Each VIOC supports up to 16 virtual interfaces; you provision the number of interfaces you need.

Figure 1-4. Logical System View shows a logical view of the system as it will be seen by the system administrator or network administrator. VNICs are at the left side of the diagram. In the center is the Zen4 system, represented as a single large device, and at the right are the physical media ports that connect to the external network. Layered behind the picture are boxes representing the hardware components that implement the various logical functions. The VNICs and a portion of the switch/router are shown as being implemented on one or more PMMs. The physical media ports and a portion of the switch are shown as being resident on one or more ENMs. The implementation of the rest of the switch is distributed across the fabric, route processor, VIOC, and ENMs. The Zen4 server is a multilevel server/switch capable of operating in Layer-2-only mode.

Figure 1-4. Logical System View

The following sections are included in Virtual I/O Controller (VIOC) and Virtualization Capabilities:

- [Virtual Network Interface Controllers](#)
- [Virtual Storage Interfaces](#)

Virtual Network Interface Controllers

Virtual network interface controllers (VNICs) represent end-points on the Layer 2 network and therefore are always assigned a Layer 2 MAC address. A user operating system (UOS) connects to the switching and routing fabric within the Zen4 system through the VNICs. The VNIC communicates with the address space of the operating system instance to

which it has been assigned, or by which it is owned. Each VNIC on a given VIOC could be owned by a different UOS instance, or several of them could be owned by the same UOS instance.

Each VNIC can be logically connected to a logical Layer-2-only port on the switch/router, as shown in .

The total bandwidth available through one VIOC is 10 Gbps, allocated to the VNICs managed by the VIOC. The bandwidth allocation is managed by VNIC rate-limiting that is implemented using a token bucket scheme. For information about configuring the VNIC bandwidth allocation, see .

Virtual Storage Interfaces

The initial virtual storage interface implementation (Zen4 Release 1.1) is as an iSCSI initiator. The iSCSI initiator interfaces to a virtual iSCSI target (typically on a storage VLAN) on the Fibre Channel Module (FCM). The iSCSI target connects to a Fibre Channel port on the FCM. The FCM Fibre Channel port is, in turn, cabled to a Fibre Channel switch or storage array.

The FCM serves as a gateway between the virtual iSCSI target and the Fibre Channel port on the FCM.

Each virtual storage interface has a nexus—a storage connection on a server, specifying the end-to-end path to the storage device. The nexus includes the virtual initiator, a virtual target interface, the world wide port name (WWPN), and the logical unit number (LUN) corresponding to the target. The nexus facility provides LUN masking and access control.

Figure 1-5. Virtual I/O Interfaces, VLANs, and Nexuses

illustrates the relationship between iSCSI initiators and targets, VLANs, and storage targets.

Figure 1-5. Virtual I/O Interfaces, VLANs, and Nexuses

Partitioning and Virtualizing Resources

Tech Company systems leverage the technology of virtualizable resources, to allow you to flexibly partition the system to match data center needs. Network interfaces, application processors, and the fabric itself can be partitioned to match the resource requirements of an application. For example, a traditional multi-tiered web-based application, as it might be deployed with traditional fixed-size boxes, requires a collection of various sized servers, routers, and appliances. As the workload expands, additional devices of the appropriate type are brought in, cabled appropriately, and deployed. As the workload shrinks, servers and routers could be uncabled and moved to a new location to be reused.

The same application can be mapped to a Zen4 system as shown in [Figure 1-6. Zen4 System Deployment in an Enterprise Data Center](#). Here the node is logically partitioned to provide the computing, network, and fabric resources demanded by the application. As the workload expands, additional partitions can be created and logically added to the application without physically moving boxes or plugging cables. As the workload shrinks, resources can be removed to be logically redeployed elsewhere, again without physically touching the boxes or the network.

Figure 1-6. Zen4 System Deployment in an Enterprise Data Center

If an application expands beyond the size of a single node, additional nodes can be added to the system. The system is still managed as a single entity, and resources from multiple nodes can be combined for a given distributed application.

shows a possible functional organization of one of the Zen4 chassis in the deployment shown in [Figure 1-7. Possible Functional Organization of a Single Zen4 System](#).

Figure 1-7. Possible Functional Organization of a Single Zen4 System

Data Flow in the Zen4 System

This section outlines the communication paths between different modules in the Zen4 chassis. The modules send and receive data traffic over these paths connected through the switch fabric. Control traffic goes over a separate management network in the system.

The following communication paths are shown in [Figure 1-8. Data Flow in the Zen4 Server](#):

1. PMM to/from PMM.
2. PMM to/from PMM (through switch fabric, IPC traffic using F7 Messaging Protocol).
3. Ethernet to/from Ethernet. (System acts like a bridge/router for this transit traffic.)
4. Ethernet to/from PMM (TCP traffic and F7MP).
5. Ethernet to/from Route Processor (through switch fabric).
6. PMM to/from local Fibre Channel ports.
7. PMM to/from route processor.
8. PMM to/from Fibre Channel Module

Figure 1-8. Data Flow in the Zen4 Server

Appendix A: Product Support

The Product Support Appendix provides information about:

- [Warranty](#)
- [Verifying Warranty Status](#)
- [Service Features](#)
- [Service Levels](#)
- [Contacting Customer Support](#)
- [Replacing a Component](#)
- [Returning a Component](#)
- [Ordering Replacement Parts or Spares](#) (for non-defective components)

Note: For additional product support information that is not covered in this appendix, visit the Tech Company Customer Support web site or call one of the customer service phone numbers listed in [Contacting Customer Support](#).

Warranty

Zen4 systems are covered by a comprehensive warranty. For detailed information about the standard warranty, contact the Tech Company Sales or Customer Support Team.

Verifying Warranty Status

Before you begin

You will need the chassis serial number from the label on the back of the Zen4 chassis.

To verify your warranty status

- Contact Tech Company Customer Support or your Tech Company account representative. (See [Contacting Customer Support](#))

Result

You will be provided all warranty status information by Tech Company for the given Zen4 chassis.

Service Features

The goal of the Tech Company Customer Support organization is to consistently deliver the quality of service you require in a mission-critical, data center environment.

- Tech Company is committed to providing around-the-clock customer service, either directly or through our service partners who have been factory-trained and certified by Tech Company Systems.
- You can log calls at any time by telephone or email, or through a web interface to our support system.

All calls are actively managed to conclusion by a Tech Company Support Engineer. (A service partner may provide the on-site repair services.) Status tracking is available via the web interface to our support system.

Service Levels

Tech Company provides multiple scalable service levels to suit individual customer needs.

Priority Levels

Some of the service levels include access to an assigned Tech Company Customer Support account engineer. This engineer can help you with installation and upgrade planning and provide progress reviews. An assigned account engineer is the primary point of contact for handling service calls.

Customers with the appropriate service contract may designate one of the following priority levels for service calls:

- Priority 1, highest severity; work starts immediately to resolve the problem.
- Priority 2, medium severity; work starts within 1 hour to resolve the problem.
- Priority 3, lowest severity; work starts within 4 hours to resolve the problem.

Contacting Customer Support

To contact Customer Support, use the phone numbers or email address in [Table A-1. Contacting Tech Company Customer Support](#).

Table A-1. Contacting Tech Company Customer Support

USA

24-hour hotline	1-800-968-6501	Email: support@Tech Company.com
-----------------	----------------	------------------------------------

Replacing a Component

This section provides the steps needed to request and receive a replacement component to Tech Company using a Return Materials Authorization (RMA).

Before You Begin

This process of requesting and replacing a component is required before you can return your defective component.

Replacing a Component

- 1. Contact the Tech Company Customer Support Center Hotline to request a replacement component (See [Contacting Customer Support](#)).**

Be prepared to provide the following information:

- Your Chassis serial number
- Serial number of the component to be returned
- Model number of the component to be returned
- Description of the problem
- Your return address
- Your phone number

- 2. Tech Company will issue you an RMA number.**

Result: You have been given an RMA number by Tech Company that represents your required authorization to return the defective component. Take note of this number, it will be needed to return your defective component.

- 3. Tech Company will ship you a replacement component.**
- 4. Unpack and install your replacement component.**
- 5. KEEP the packaging that the replacement component arrived in.**

Result

You have received an authorization to return your defective component, an RMA number, and your replacement component.

The RMA number and the replacement component packaging will be used to return your defective component.

Returning a Component

This section provides the steps needed to return a defective component to Tech Company after receiving a replacement component (See [Replacing a Component](#)).

Before You Begin

To return a defective component, you must first request and receive a replacement component.

You will need your Return Materials Authorization (RMA) number and the packaging that your replacement component arrived in to complete this process.

Returning a Defective Component

- 1. After you have unpacked and installed the replacement component, use the same packaging materials to pack the defective item for return to Tech Company.**
- 2. Clearly mark the RMA number on the packaging exterior.**

NOTE: Only packages with RMA numbers written on the outside of the shipping carton and/or the packing slips and shipping paperwork will be accepted by the Tech Company Receiving Department. All other packages will be rejected.

- 3. Ship the product back to Tech Company using the carrier designated by Tech Company.**

The carrier information will be provided to you by Tech Company.

NOTE: Tech Company will be responsible for all freight charges for returned products or components as long as you use the carrier designated by Tech Company.

- 4. Tech Company receives and processes the returned component.**

The title to any returned products or components will transfer to Tech Company upon receipt.

Result

You have returned your defective component and the title of the component has been transferred to Tech Company.

Ordering Replacement Parts or Spares

Before you begin

Replacement is for a non-defective component. See [Replacing a Component](#) and [Returning a Component](#) for the steps to replace a defective component.

Ordering replacements parts or spares

- Visit the Tech Company Customer Support web site or contact your Tech Company account representative to discuss ordering a replacement component or spare part. (See [Contacting Customer Support](#))

Result

Tech Company will assist you in ordering a replacement component or spare part.

Glossary

M

My Term
My definition

Index

C

clock, system 17
customer support 23

E

ENM
 overview 17
Ethernet Network Module. See ENM 17

F

fan tray
 redundant 18
flash, on-board 17

H

HyperTransport links 17

I

iSCSI initiators 19

M

modules
 ENM 17

O

Opteron 800 series processor 17
overview, system 6

P

physical resources 16
power
 redundant 18
Processor and Memory Module 16
product support 23

Q

Quaternia system module 17

R

redundant subsystems 18
Reliability, availability, and serviceability (RAS) 9
replacements, ordering 31
replica, of critical data structures 17
resources 9

S

SMP 17
spares, ordering 31
subsystem failover 17
subsystems, redundant 18
symmetric multiprocessing 17
system
 clock 17
system overview 6

V

VNICs 19

Z

Zen4 subsystems 17
ZSM 17