NonStop Himalaya Range: K200, K2000, and K20000 Servers

Key features

System architecture designed for critical applications

- Power of parallelism
- Scalable servers
- Unmatched reliability

NonStop Kernel

- Fault tolerance and data integrity
- Parallel processing software
- Single-system image
- Kernel personalities

Distributed processing and distributed data

- Designed for highly distributed environments
- Cost-effective manageability
- Reliability in a distributed environment

Multivendor network and ATM integration

Open client/server computing

A choice of application development tools

Investment protection

Automated diagnostic system

Fast, efficient, online servicing

Himalaya range: from entry-level to massively parallel servers

Specifications

Open, reliable, parallel servers offering excellent price/performance

The NonStop Himalaya range of servers combines Tandem's unique parallel processing architecture with industry-standard RISC technology to deliver outstanding price/performance for online transaction processing (OLTP), client/server, messaging, and decision support applications.

The Himalaya servers and Tandem's open NonStop Kernel together deliver continuous availability, data integrity, online growth with linear scalability, networking, security, distributed processing, and distributed data capability.

TANDEM

Himalaya servers provide a wide performance range. The entry-level K200 server delivers high reliability for LAN and server environments. To meet an even wider variety of requirements, the K2000 server offers higher levels of scalability and connectivity. The massively parallel K20000 server is the most powerful and most reliable commercial server ever built.

With scalability from 2 to more than 4,000 processors, Himalaya servers can run your largest critical business applications—centralized or distributed—at peak performance. Each server appears as a single entity to applications and users, regardless of the physical distribution of the nodes. Growth within a server and on the network can occur online without disrupting running applications.

To protect your investment in existing solutions, Himalaya servers are fully application compatible with all NonStop servers. You can move applications among platforms with little or no modification.

To minimize operating costs, most components of the Himalaya servers are designed to be installed and serviced by users. In addition, diagnostics can be run from a remote service center without disrupting system operation.

For maximum configuration flexibility, Himalaya servers—whose compact system design minimizes floorspace requirements—can operate in a non-computer-room environment.

System architecture designed for critical applications

Using advanced workstation technology and Tandem's parallel processing architecture, the Himalaya range of servers delivers outstanding price/performance across a wide variety of commercial applications.

Power of parallelism

To meet the demanding requirements of today's applications, parallelism is designed throughout the hardware and software architecture of Himalaya servers: in its processors, communications controllers, I/O channels, storage devices, the NonStop Kernel, and application processes. Parallelism improves performance by transparently distributing transactions or queries across multiple processors. The parallel architecture of Himalaya servers and the message-based NonStop Kernel provide easy scalability within a single server or across a network of Tandem servers.

Scalable servers

The Himalaya architecture is composed of cost-effective modular units that facilitate incremental processor and I/O growth. Key modules are designed for insertion and removal—while the server is running—without data corruption or application failure.

The K20000 server, for example, can be scaled from 2 to more than 4,000 processors, providing practically unlimited growth potential without any application changes. Tandem's TorusNet technology provides high-speed interconnectivity in these massive configurations. TorusNet sections, composed of four processing modules, are connected horizontally and vertically to form domains of up to 224 processors. TorusNet domains can be interconnected to provide the parallel processing strength of more than 4,000

processors, enough to handle today's most demanding commercial transaction, messaging, and information access applications (see figure 1).

As your applications and networks expand, Tandem's unique loosely coupled architecture delivers linear performance improvement: as each processor is added, a full processor's worth of performance is delivered to the application.

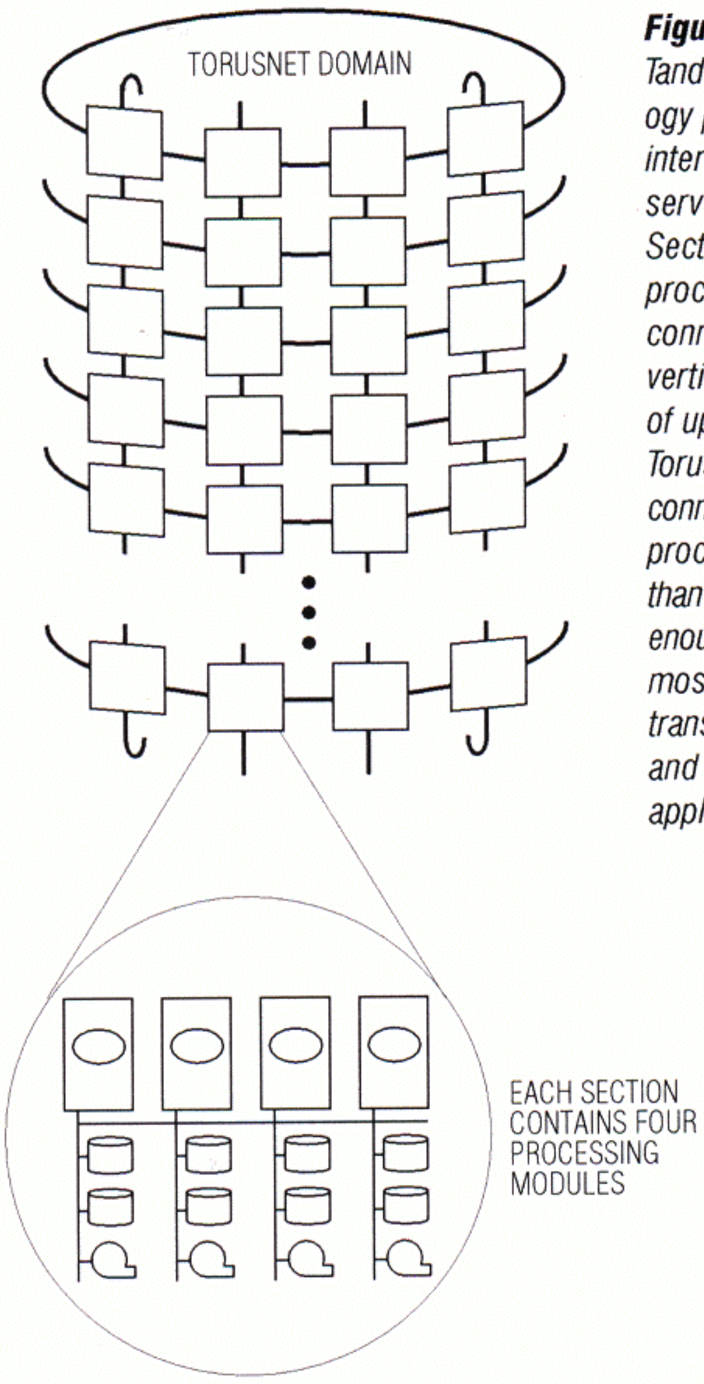

Figure 1
Tandem's TorusNet technology provides high-speed interconnectivity of massive server configurations. Sections, composed of four processing modules, are connected horizontally and vertically to form domains of up to 224 processors. TorusNet domains can interconnect to provide the parallel processing strength of more than 4,000 processors, enough to handle today's most rigorous commercial transaction, messaging, and information access applications.

Unmatched reliability

The reliability of Himalaya servers is unmatched in the industry, providing continuous availability for your applications and data integrity for your crucial transactions.

The loosely coupled architecture of Himalaya servers consists of multiple processors, dual interprocessor buses, dual-ported controllers, and fault-tolerant power subsystems. The architecture prevents all single and most multiple hardware or software malfunctions from disrupting your applications. In this parallel processing architecture, the workload is shared among processors that perform multiple tasks simultaneously. Under normal operation, all processors share the workload; there are no idle backup units. In the event that one processor fails, its workload is automatically shared by the other processors. Should there be a power outage, server memory is preserved via integrated battery backup modules.

The design of Himalaya servers includes many features that provide data integrity, including parity and control sequence checking on buses, error-correcting codes (ECC), and battery backup on memory, as well as end-to-end sector and block checksums. A processor is made up of two microprocessors running in lockstep and executing the same code. The output of each microprocessor is constantly compared to that of the other. If the lockstepped microprocessors disagree, the processor immediately shuts down, thus preventing any corruption of data. At that point, the NonStop Kernel transparently transfers the workload to other processors, providing continuous service to applications.

NonStop Kernel

The Himalaya range is based on the Tandem NonStop Kernel, designed for massive scalability. With the NonStop Kernel, virtually all limits to database

size and transaction throughput have been removed. In addition, the NonStop Kernel is the core of Tandem's open systems environment. Together, the Himalaya servers and NonStop Kernel deliver unprecedented reliability, parallelism, and price/performance in an open computing environment.

Fault tolerance and data integrity

The NonStop Kernel is designed for high reliability, including fault tolerance and data integrity features.

Because Himalaya servers use loosely coupled processors, a software failure in one processor does not disrupt processing: the workload is taken over by the other processors.

The NonStop Kernel also contributes to continuous availability by enabling you to run primary and backup processes in separate processors for all applications. Before performing any critical function, the primary process sends the backup process a checkpoint message. The message contains the data and status information the backup process needs to complete the operation should a problem occur in the primary process. Checkpointing is streamlined, using only a minimal amount of the processor's time. This software fault tolerance enables Tandem NonStop servers to deliver 10 times the availability of other vendors' systems.

The NonStop Kernel also supports end-to-end data integrity features, ensuring delivery of critical business transactions.

Parallel processing software

Using message-based microkernel technology, the NonStop Kernel allows processing power to be used in parallel for optimum performance. The NonStop Kernel also allows database, transaction, and message services to be executed in parallel to increase overall throughput and reduce response time.

Single-system image

The NonStop Kernel is the industry's only software that provides a single-system image across all servers of a distributed network. The single view dramatically improves the ability to deploy and manage highly parallel applications, reducing operations costs. In addition, the NonStop Kernel supports Tandem's TorusNet technology for high-speed data transmission in a manageable, reliable, and scalable application environment that supports thousands of processors.

Kernel personalities

The NonStop Kernel supports application program and operations control interfaces, called "personalities," that allow a diverse set of application software to run on Himalaya servers and derive the benefits of scalability, reliability, and outstanding price/performance (see figure 2). Providing the foundation for these personalities, both Guardian Services (for compatibility with existing OLTP applications) and Open System Services (consisting of widely used UNIX® interfaces) are integrated on the server, providing maximum flexibility.

In addition, message, transaction, and database services are integrated with the NonStop Kernel. With relational database management personalities, for example, applications can be developed using SQL, Data Access Language (Macintosh), SQL Server (Microsoft/Sybase), ODBC (Microsoft), or Oracle Tools (Oracle). Transaction processing personalities include Parallel Transaction Processing (PTP) Services for the CICS™ API and NonStop TUXEDO®. PTP Services is compatible with applications for IBM's CICS™ products.

Distributed processing and distributed data

Tandem's system architecture and the NonStop Kernel environment provide many features that allow you to choose where and how to distribute your applications and data to maximize reliability and cost-effectiveness.

Designed for highly distributed environments

Within Tandem's system architecture, database, transaction, and messaging services have all been designed to support highly distributed but easily managed application environments. And we designed the application environment to support a parallel, loosely coupled architecture.

For example, NonStop SQL/MP, Tandem's implementation of the ANSI-standard structured query language, enables programmers to design applications to manage the most complex distributed database. Data can reside on a single Himalaya server or be distributed on multiple Himalaya servers throughout the network. Although data is distributed to various locations in the network, a distributed database appears to users and applications as a single database.

Cost-effective manageability

Tandem's NonStop NET/MASTER and Distributed Systems Management (DSM) services and applications enable you to monitor and control computer resources efficiently. Applications, users, and administrators can access resources and information anywhere within a distributed network as though they were local resources.

DSM provides a single management system for collecting, logging, and distributing management data from a single server or a network of Himalaya servers.

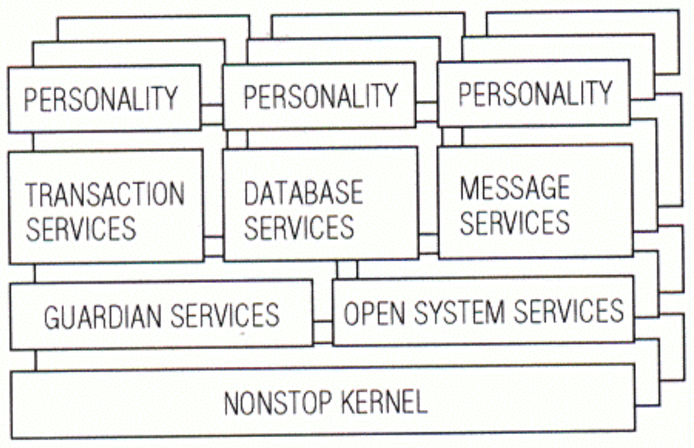

Figure 2
The NonStop Kernel supports application program interfaces, called personalities, that allow a diverse set of applications to run on NonStop servers and transparently derive the advantages of scalability, reliability, and outstanding price/performance. The NonStop Kernel supports both Guardian Services (for compatibility with existing Tandem applications) and Open System Services (consisting of widely used UNIX interfaces), which are integrated on the server, providing maximum flexibility.

This allows you to effectively manage highly distributed systems, no matter how massive or geographically dispersed.

Reliability in a distributed environment

Tandem was the first vendor to introduce a fully distributed, update-protected database founded on open standards; we pioneered the concept of two-phase commit in a network environment. As a result, you can benefit from our experience in working with businesses to implement and support highly distributed solutions that depend on the reliability of our software and hardware.

The data integrity and availability features of the Himalaya range continue to provide the highest levels of reliability—however or wherever you choose to distribute applications, data, or servers.

The NonStop Kernel provides data integrity within a single NonStop server or worldwide network of servers.

Tandem's NonStop Transaction Manager/MP (NonStopTM/MP) protects your data from the effects of incomplete transactions, system failures, or network failures. To ensure transaction integrity, NonStop TM/MP attempts to complete the transaction by updating all other relevant servers in the network; if that is not possible, it backs out the transaction.

In the event of an environmental disaster, Tandem's Remote Duplicate Database Facility (RDF) provides rapid recovery of

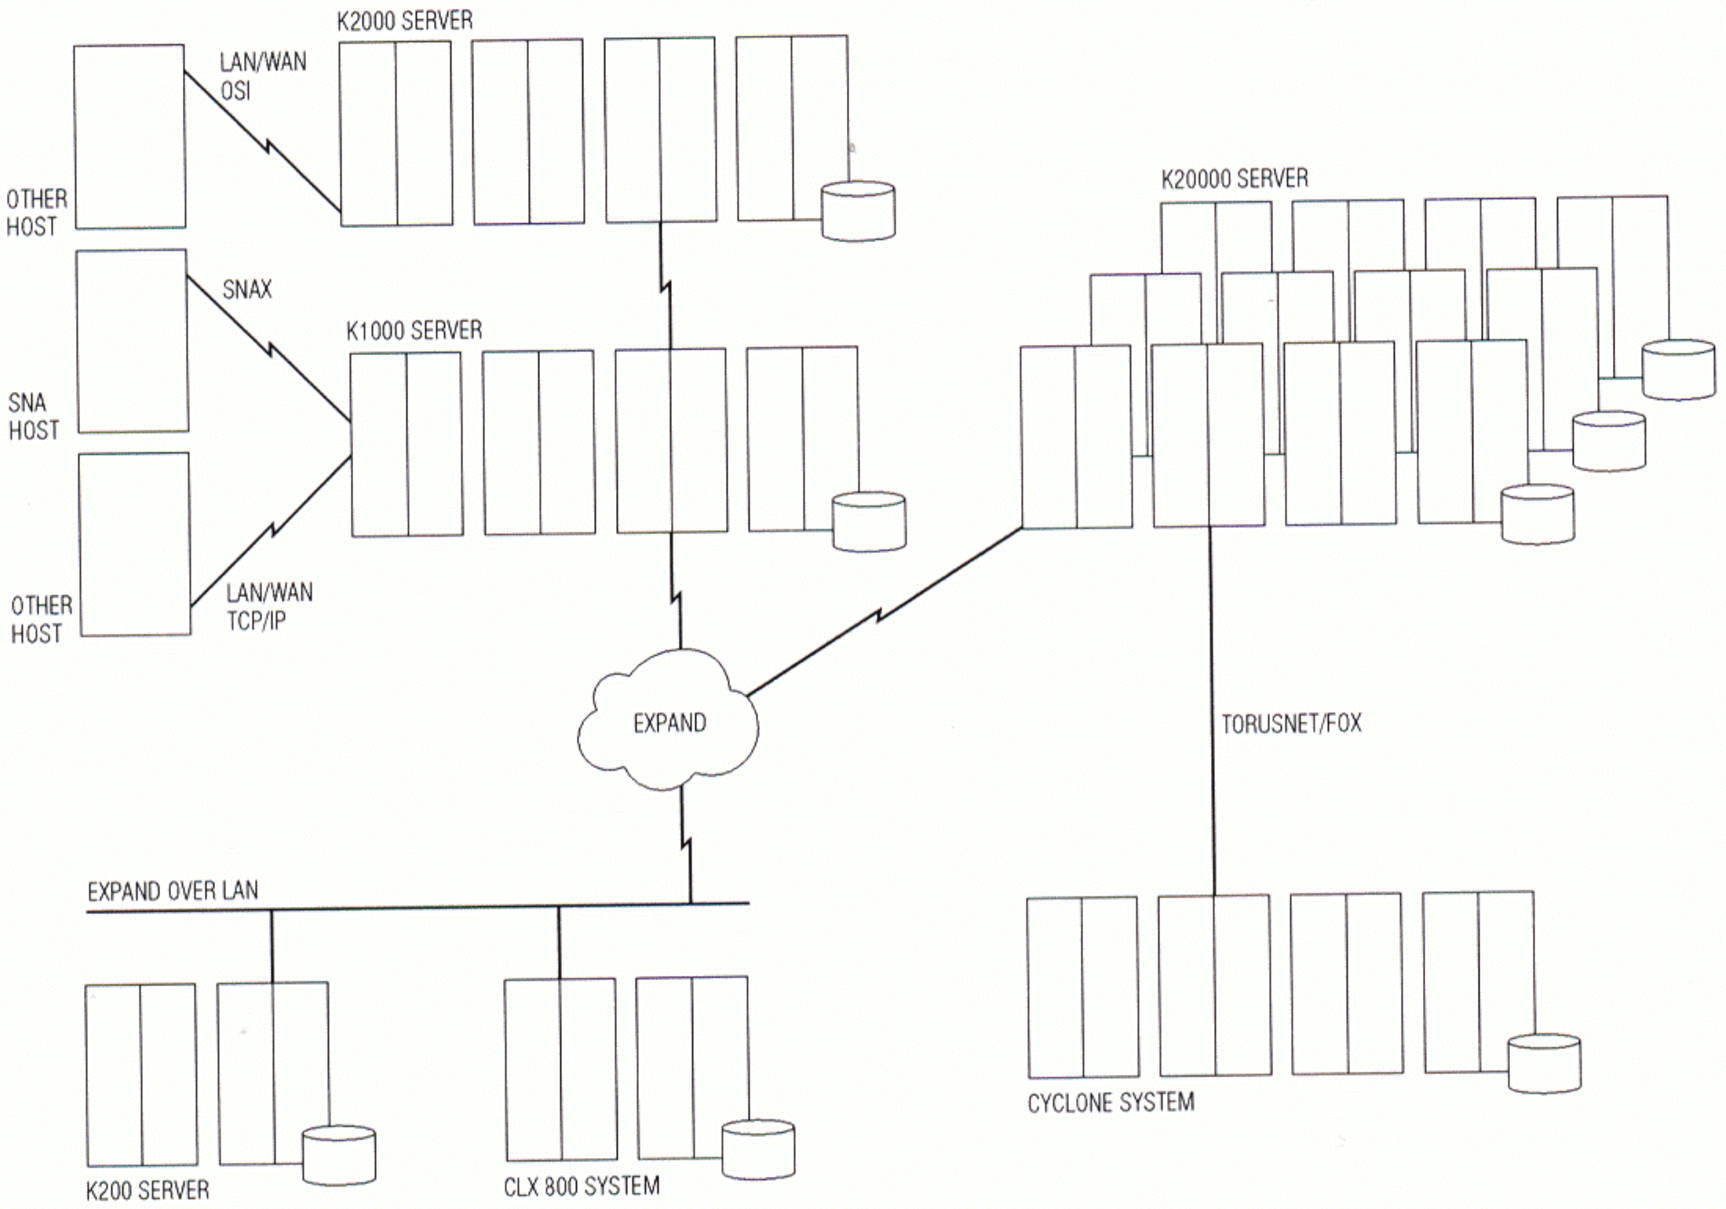

Figure 3
NonStop Himalaya servers support a wide range of open networking standards, including TCP/IP, NETBIOS, AppleTalk, SNA, and OSI. In addition, Himalaya servers support protocols that enable you to integrate various standard and non-standard devices.

your database by letting you maintain a current online copy of all or part of the database at another site. The presence of the backup is transparent to applications.

If you are running NonStop TM/MP, no application code has to be changed to implement RDF. In addition, RDF allows users to perform read operations from remote duplicate databases, improving response times.

To protect your data from unauthorized use, unauthorized modification, or destruction, Tandem's Safeguard security management facility enables you to implement a full range of organizational security policies for your NonStop servers and their users.

Tandem is certified at the C2 level of the National Computer Security Center (NCSC) security standards and was the first vendor to be certified to ITSEC F-AV/E3 Availability Standards, as supported by the European Commission.

Multivendor network and ATM integration

Himalaya servers support a wide range of open networking standards, including TCP/IP, IPX/SPX, NETBIOS, AppleTalk, SNA, and OSI. In addition, they support protocols that enable you to integrate a wide range of standard and nonstandard devices into your information system (see figure 3).

Asynchronous Transfer Mode (ATM) communications are provided using the Himalaya ATM Starter Kit. This kit consists of UB Networks' GeoRim/E product and Tandem's 3615 High-Performance Ethernet controller.

Ethernet connectivity is provided by Tandem's high-performance LAN controller, designed for the demands of peak-load, low-latency, client/server use. NonStop Himalaya servers can interface to a variety of LAN-based devices including PCs (running Microsoft DOS and Windows), UNIX workstations, and Apple Macintosh computers.

Open client/server computing

NonStop Himalaya servers are ideal for open client/server applications. Tandem's client/server products allow applications to take advantage of the graphical user interfaces and low-cost processing power of workstations while gaining the speed and transaction protection of Tandem servers.

Himalaya servers can connect to a wide range of client platforms, including Microsoft Windows, native MS-DOS, Presentation Manager, native OS/2, SCO UNIX, Sun SparcStation SunOS, HP 9000 series 700 systems, Microsoft NT for Intel, and Macintosh. Himalaya servers also support a wide range of open database gateways.

A choice of application development tools

Tandem's open application development environment can make your programmers more productive. To facilitate the design and development of your commercial applications, Himalaya servers support industry-standard programming languages and leading brand-name tools for all aspects of client/server application development.

Investment protection

Applications developed and deployed on other NonStop systems will run on Himalaya servers, with little or no modification. To protect your investment in I/O devices and controllers, the Himalaya servers connect to other NonStop system I/O subsystems. Most controllers and devices that attach to NonStop Himalaya K10000, K1000, K100, Cyclone, Cyclone/R, CLX, and CLX/R systems are supported by Himalaya K20000, K2000, and K200 servers.

Automated diagnostic system

The maintenance and diagnostic system of Himalaya servers automatically monitors the status of processors, I/O controllers, power subsystems, and cabinet environments. Key data is logged and evaluated using expert-system techniques to identify potential problems. Monitoring and evaluation take place while the system is online, keeping applications up and available to users. Using the Remote Maintenance Interface (RMI) hardware and software, Himalaya servers can dial out to an online support center to report potential problems before they impact operations.

Diagnostics can be run from a remote service center without disrupting server operation. When a problem is identified, corrective action can be initiated from the service center.

Fast, efficient, online servicing

Most components of Himalaya servers—including processors, I/O controller boards, storage devices, power supplies, and fans—can be installed and serviced by users. Virtually all of the field-replaceable units (FRUs) can be installed and serviced online without disrupting the normal operation of the server.

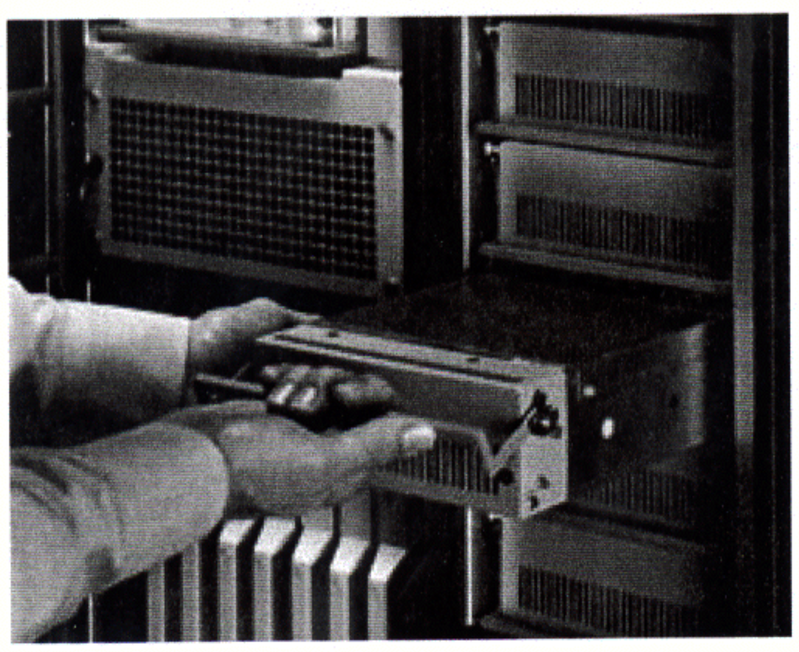

Most components of Himalaya servers—including processors, I/O controller boards, storage devices, power supplies, and fans—are designed to be installed and serviced by users.

To meet your specific needs, Tandem offers a variety of service and support options. These programs are complemented by system features that support remote access for diagnostics and repair—to maximize system availability and minimize operating costs.

Himalaya range: from entry-level to massively parallel servers

Himalaya K200 entry-level parallel servers

Himalaya K200 servers are expandable, entry-level parallel servers based on reduced instruction-set computing (RISC) microprocessor technology.

The K200 RISC processor comes with up to 128 megabytes of onboard error-correcting memory, 1 megabyte of secondary cache, and dual interprocessor buses. From its smallest to largest configuration, the Himalaya K200 server offers a twofold performance growth range.

K200 servers come standard with two multifunction controllers in each processor cabinet, and slots for one tape drive and up to six disk drives. The servers use low-cost internal 4255 disk drives that store 2 gigabytes of formatted data per disk drive and fit inside the K200 processor cabinet.

Himalaya K200 servers: packaged configurations

	K202	*K204*
Processor cabinets	1	2
Processors (R4400, 125 MHz)	2	4
Memory (MB)	128	256
Multifunction controllers	2	4
Ethernet controllers	1	1
Unassigned I/O slots	2	5
Communications lines		
Asynchronous	4	8
Synchronous	2	4

The K200 servers use direct memory access (DMA) I/O channels, so processor cycles are not required to move data from a channel to memory or vice versa. This feature is especially important for I/O-intensive applications.

Himalaya K2000 scalable parallel servers

The Himalaya K2000 servers provide a cost-effective platform for applications requiring midrange performance and a wide range of scalability. Using advanced RISC microprocessor technology and Tandem's parallel system architecture, the K2000 servers deliver all of the Tandem advantages—including high performance, reliability, and scalability—at outstanding price/performance levels.

Each K2000 processor comes with 128 megabytes of onboard error-correcting memory expandable to 256 megabytes, high-speed cache memory of 1 megabyte, and dual interprocessor buses. A processor cabinet also comes standard with two multifunction controllers and slots for six disk devices. For I/O expansion, the K2000 server provides a single I/O channel per processor and allows attachment of one multifunction I/O cabinet (per processor cabinet), each with eight I/O controller slots and six disk slots.

The K2000 servers use direct memory access (DMA) I/O channels, so processor cycles are not required to move data from a channel to memory or vice versa.

Himalaya K20000 massively parallel commercial servers

The K20000 massively parallel server is the most powerful commercial server ever built. With scalability from 2 to more than 4,000 processors, it lets you run your largest applications at peak performance.

The Himalaya K20000 processor is the third-generation RISC microprocessor for NonStop servers. It is based on the MIPS® R4400™ microprocessor clocked at 200 megahertz, so it leverages the economics of workstation technology while providing the Tandem advantages of data integrity and continuous availability —key requirements in high-volume commercial applications.

Each K20000 processor has 128 megabytes of onboard error-correcting memory expandable to 256 megabytes, high-speed secondary cache memory of 4 megabytes, and dual interprocessor buses.

The K20000 server was designed to accommodate the most demanding I/O requirements. Each K20000 processor comes standard with two I/O channels. Two additional I/O channels can be added to each processor to provide a total of four channels with an aggregate I/O bandwidth of 20 megabytes per second per processor. A fully configured 16-processor K20000 server has 64 I/O channels. Each channel supports up to 32 I/O controllers. These controllers work with a variety of devices and communications lines, enabling your K20000 server to support thousands of gigabytes of disk storage and hundreds of thousands of clients.

The K20000 servers use direct memory access (DMA) I/O channels, so processor cycles are not required to move data from a channel to memory or vice versa.

Himalaya K2000 servers: packaged configurations

	K2002	*K2004*	*K2006*	*K2008*	*K2012*	*K2016*
Processors (R4400, 125 MHz)	2	4	6	8	12	16
Memory (MB)	256	512	768	1,024	1,536	2,048
Multifunction controllers	2	4	6	8	12	16
Ethernet controllers	1	1	1	1	1	1
Unassigned I/O slots	3	7	11	15	23	31
Communications lines						
Synchronous	2	4	6	8	12	16
Asynchronous	4	8	12	16	24	32

Himalaya K20000 servers: packaged configurations

	K20002	*K20004*	*K20006*	*K20008*	*K20012*	*K20016*
Processors (R4400, 200 MHz)	2	4	6	8	12	16
Memory (MB)	256	512	768	1,024	1,536	2,048
Multifunction I/O cabinets	1	1	1	1	1	1
Ethernet controllers	1	1	1	1	1	1
Unassigned I/O slots	7	7	7	7	7	7
Communications lines						
Asynchronous	4	4	4	4	4	4
Synchronous	2	2	2	2	2	2
TorusNet H-link controllers	—	—	4	4	6	8

Each processor cabinet contains 2 processors and can be attached to up to 4 multifunction I/O cabinets, each with 8 I/O slots (MF-8), for office environment operation; or 4 multichannel I/O cabinets, each with 32 slots (MC-32), for massive corporate configurations. The K20000 packaged configurations come standard with one multifunction I/O cabinet.

Himalaya processor specifications

Configurations

	K200	*K2000*	*K20000 TorusNet Node*	*K20000 TorusNet Domain*	*K20000 Multidomain TorusNet*
Micro-processor type	MIPS R4400/ 125 MHz	MIPS R4400/ 125 MHz	MIPS R4400/ 200 MHz	MIPS R4400/ 200 MHz	MIPS R4400/ 200 MHz
Maximum processors	4	16	16	224	4,080
Relative performance range[1]	2 to 4	2 to 16	3.9 to 31.5	3.9 to 440.9	3.9 to 8,030.5
Cache (per processor)	1 MB	1 MB	4 MB	4 MB	4 MB
Maximum main memory[2]	512 MB	4,096 MB	4,096 MB	57 GB	1,044 GB
Maximum disk storage[3]	184 GB	3,232 GB	65 TB	918 TB	16,711 TB
Maximum I/O channels	4	16	64	896	16,320
Multifunction controllers	2/proc. cabinet	2/proc. cabinet	—	—	—
Minimum battery hold-up time	3.75 hr	2 hr	2 hr	2 hr	2 hr
NonStop Kernel	Release D30.01 or later	Release D30.01 or later	Release D30.01 or later	Release D30.01 or later	Release D30.01 or later

Operating environment

	K200	*K2000*	*K20000*
Temperature			
Operating	5° to 40° C (41° to 104° F)	5° to 40° C (41° to 104° F)	5° to 40° C (41° to 104° F)
Nonoperating	–29° to 55° C (–20° to 131° F)	–29° to 55° C (–20° to 131° F)	–29° to 55° C (–20° to 131° F)
Relative humidity			
Operating	10% to 85%, noncondensing	10% to 85%, noncondensing	10% to 85%, noncondensing
Nonoperating	10% to 95%, noncondensing	10% to 95%, noncondensing	10% to 95%, noncondensing
Power dissipation (full cabinet)	1,380 W	1,400 W	1,520 W

[1] Performance range is stated relative to the K200 entry-level server (K202=2 units).

[2] Main memory is ECC, providing single-bit error correction, double-bit error detection.

[3] Measurement is based on a single controller per string.

Operating environment *(continued)*

	K200	*K2000*	*K20000*
Altitude			
Operating	3,048 m (10,000 ft)	3,048 m (10,000 ft)	3,048 m (10,000 ft)
Nonoperating	15,240 m (50,000 ft)	15,240 m (50,000 ft)	15,240 m (50,000 ft)
Heat dissipation	1,580 W (4,712 Btu/hr) per cabinet	1,600 W (4,780 Btu/hr) per cabinet	2,105 W (5,190 Btu/hr) per cabinet
Noise	58 dBa	58 dBa	58 dBa

Power *(all processor models)*

100–120 volt model	
Nominal rating	100–120 V AC, single phase, 50–60 Hz
Operating range	88–132 V AC, continuous, 48–62 Hz
Line current, 2 cords	12 A RMS, maximum, per cord
Power cord/plug	Two 8-ft cords, 20 A plug (NEMA 5-20P)
200–240 volt model	
Nominal rating	200–240 V AC, single phase, 50–60 Hz
Operating range	177–264 V AC, continuous, 48–62 Hz
Line current, 2 cords	6 A RMS, maximum, per cord
Power cord/plug	Two 2.5-meter detachable power cords. Choice of: NEMA 5-20A (USA, Canada, Japan); NEMA 6-20A (W. Japan); CEE (7) VII type VIIG (Continental Europe); AS 3112 type SAA/3 (Australia, New Zealand); BS 1363/A type BS/89/13A (UK, Ireland); IEC 309 type 248 (Switzerland, Denmark); CEI 23-16 type I/3/16 (Italy); SI 32/1971 type IL/3 (Israel); CEE (7) VII type VIIG (4) (Brazil, Argentina); SABS/164/1980 type SABS/16 (South Africa)

Physical

	K200	*K2000*	*K20000*
Height	105.79 cm (41.65 in)	140 cm (55 in)	140 cm (55 in)
Width	53.97 cm (21.25 in)	48 cm (19 in)	48 cm (19 in)
Depth	69.21 cm (27.25 in)	76 cm (30 in)	78 cm (30.5 in)
Weight (maximum)	205 kg (450 lb)	264 kg (580 lb)	264 kg (580 lb)
Floorspace requirements			
Cabinet only	0.37 sq m (4.0 sq ft)	0.37 sq m (4.0 sq ft)	0.37 sq m (4.0 sq ft)
Front clearance	61 cm (24 in)	61 cm (24 in)	61 cm (24 in)
Back clearance	76.2 cm (30 in)	76.2 cm (30 in)	76.2 cm (30 in)

Multifunction I/O cabinet (MF-8) subsystems specifications

Configurations

	K200	*K2000*	*K20000*
Maximum cabinets per processor pair	—	1	4
Maximum cabinets per 16-processor node	—	8	32
Multifunction controllers per cabinet	—	2	2
Disk/tape slots per cabinet	—	6	6
I/O controller slots per cabinet	—	8	8

Operating environment

Temperature	
Operating	5° to 40° C (41° to 104° F)
Nonoperating	–29° to 55° C (–20° to 131° F)
Relative humidity	
Operating	10% to 85%, noncondensing
Nonoperating	10% to 95%, noncondensing
Power dissipation (full cabinet)	1,600 W
Altitude	
Operating	3,048 m (10,000 ft)
Nonoperating	15,240 m (50,000 ft)
Heat dissipation	1,750 W (5,463 Btu/hr)
Noise	58 dBa

Power

100–120 volt model	
Nominal rating	100–120 VAC, single phase, 50–60 Hz
Operating range	88–132 VAC, continuous, 48–62 Hz
Line current, 2 cords	16 A RMS, maximum, per cord
Power cord/plug	Two 8-ft cords; 20 A plug (NEMA 5-20P)
200–240 volt model	
Nominal rating	200–240 V AC, single phase, 50–60 Hz
Operating range	177–264 V AC, continuous, 48–62 Hz
Line current, 2 cords	8 A RMS, maximum, per cord
Power cord/plug	Two 2.5-meter detachable power cords. Choice of: NEMA 5-20A (USA, Canada, Japan); NEMA 6-20A (W. Japan); CEE (7) VII type VIIG (Continental Europe); AS 3112 type SAA/3 (Australia, New Zealand); BS 1363/A type BS/89/13A (UK, Ireland); IEC 309 type 248 (Switzerland, Denmark); CEI 23-16 type I/3/16 (Italy); SI 32/1971 type IL/3 (Israel); CEE (7) VII type VIIG (4) (Brazil, Argentina); SABS/164/1980 type SABS/16 (South Africa)

Physical

Height	140 cm (55 in)
Width	49 cm (19 in)
Depth	78 cm (30.5 in)
Weight (maximum)	259 kg (570 lb)
Floorspace requirements	
Cabinet only	0.37 sq m (4.0 sq ft)
Front clearance	61 cm (24 in)
Back clearance	76.2 cm (30 in)

Multichannel I/O cabinet (MC-32) subsystems specifications

Configurations

	K200	*K2000*	*K20000*
Maximum cabinets per processor pair	—	—	4
Maximum cabinets per 16-processor node	—	—	32
I/O controller slots per cabinet	—	—	32

Operating environment

Temperature	
Operating	16° to 30° C (60° to 86° F)
Nonoperating	–29° to 55° C (–20° to 131° F)
Relative humidity	
Operating	42% to 70%, noncondensing
Nonoperating	10% to 95%, noncondensing
Power dissipation (full cabinet)	3,570 W
Altitude	
Operating	3,048 m (10,000 ft)
Nonoperating	15,240 m (50,000 ft)
Heat dissipation	3,570 W (12,189 Btu/hr)
Noise (full cabinet)	58 dBa

Power[4,5]

U.S. model—operating range, normal operation	
Voltage	177–225 V AC, 3 phase, continuous
Frequency	58–62 Hz
Phase current	14 A RMS, maximum
Neutral current	14 A RMS, maximum
Line cord	8 ft
Line cord plug	IEC 309, 530PG
International model—operating range, normal operation	
Voltage	325–456 V AC, 3 phase, continuous
Frequency	58–62 Hz
Phase current	10 A RMS, maximum
Neutral current	10 A RMS, maximum
Line cord	8 ft
Line cord plug	Not included

Physical

Height	185 cm (73 in)
Width	91 cm (35.75 in)
Depth	85 cm (33.5 in)
Weight (maximum)	522 kg (1,150 lb)
Floorspace requirements	
Cabinet only	0.78 sq m (8.3 sq ft)
Front clearance	92 cm (36 in)
Back clearance	92 cm (36 in)

[4] The multichannel I/O cabinet power system assumes that all system cabinets, including disk and tape cabinets, are powered from the same site power distribution panel.

[5] AC power receptacles are not shipped with the multichannel I/O cabinetry. They must comply with local regulations and power requirements.

For more information
For additional information, please contact your Tandem representative or access Tandem's home page on the World Wide Web at http://www.tandem.com/. Send Internet e-mail messages to info@tandem.com.

Tandem Computers Incorporated

World Headquarters
19333 Vallco Parkway
Cupertino, CA 95014-2599, USA
+1 (408) 285 6000

Americas Division Headquarters
19191 Vallco Parkway
Cupertino, CA 95014-2594, USA
USA: +1 (800) 482 6336
Canada: +1 (800) 345 8636
Latin America: +1 (408) 285 3738

European Division Headquarters
Antareslaan 11
2132 JE Hoofddorp, Netherlands
+31 (2503) 68 000

Asia-Pacific Division Headquarters
300 Beach Road #33-01
The Concourse, Singapore 0719
+65 297 4866

Japan Division Headquarters
Tennoz Central Tower 2-2-24
Higashi-Shinagawa, Shinagawa-ku
Tokyo 140, Japan
+81 (3) 5463 6600

Offices in the United States, Australia, Brazil, Canada, Chile, Europe, Hong Kong, Japan, Mexico, New Zealand, People's Republic of China, and Singapore.

Distributors and other agents in Argentina, Colombia, Finland, Greece, Guatemala, India, Indonesia, Israel, Korea, Malaysia, the Middle East, Peru, the Philippines, Portugal, Taiwan, Thailand, Turkey, Uruguay, and Venezuela.

CD1095-0195

TANDEM