Modular 1U Rackmount Server Design Enhances Data Center Flexibility

A Trenton Systems White Paper

Contents

TRENTON SYSTEMS®
Engineered For Reliability

Introduction

The term "modular" is one of those over-used and often times misused terms that engineers like to throw around. Let's reset our thinking by defining the term modular. According to Dictionary.com modular means, "*composed of software or hardware modules that can be altered or replaced without affecting the remainder of the system*," while the term "modules" is defined as, "*a separable component, frequently one that is interchangeable with others, for assembly into units of differing size, complexity, or function*." The implication of the term modular when it comes to a hardware computing platform is one of simplicity. Of course we know that designing in simplicity can be a complex engineering task.

Many modular computer designs tend to focus on large, expensive, purpose-built computer chassis. Wouldn't it be great if we could apply that same high-end modular computer design approach in order to construct a general purpose 1U rackmount computer hardware platform? If we could do this we could have a flexible 1U platform that could be deployed in a wide variety of data center system architectures. In the following paper, we'll talk about how to accomplish this while utilizing a modular 1U rackmount computer design that uses embedded components to maximize system longevity and by default software architecture stability. Modularity also ensures easy upgradability and field support.

Section 1 - Rackmount Server Modularity Overview

Many of today's 1U rackmount servers are designed for short-term data center applications. Modularity and long-term hardware platform stability requirements are frequently ignored. The modular design approach illustrated in Figure 1 below is very different from the typical 1U rackmount server.

The chassis covers have been made transparent in the Figure 1 CAD drawing to provide a quick overview of the server's modular chassis design concepts. The modular processor board or blade is mounted in a metal carrier to provide:

1) Processor board stability
2) An I/O and network interface bracket
3) Fast and easy insertion and removal from the 1U rackmount server.
4) Secure mounting and access for a full-length, industry standard, PCI Express plug-in edge card

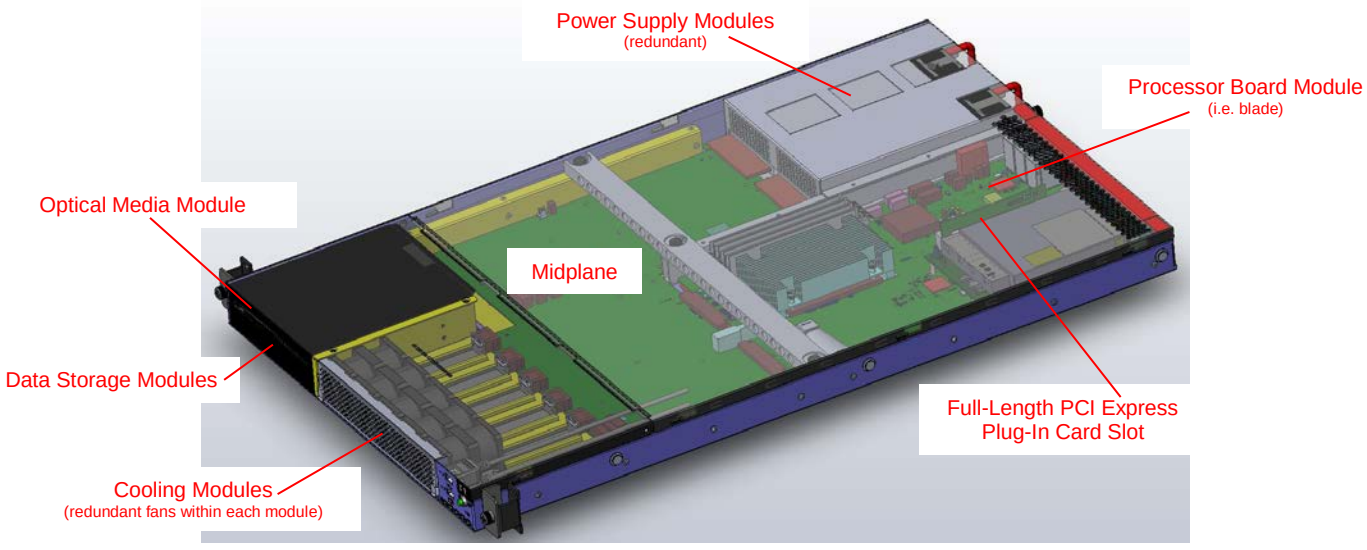

Figure 1 – MBS1000 1U Rackmount Server CAD Drawing – Top Chassis Covers Transparent

The modularity design concepts built into this 1U rackmount server; that we will call the MBS1000, extends to other rackmount server components. While the chassis design features a passively cooled processor blade, the chassis itself has five (5) modular cooling fan assemblies. Each cooling fan module has two fans. One fan is active while the other fan on the module is for backup. The MBS1000 is equipped with Smart System Management (SSM) software that is loaded into the processor board's Baseboard Management Controller (BMC) at the time of manufacture. This software monitors fan speed and temperature to provide fan speed control and to switch over to the secondary backup fan should the first fan exhibit signs of an impending failure. If this were to happen, status messages are transmitted to a central location via the chassis' dedicated system management LAN. A technician can pop off the chassis filter cover, slide out the fan module with the problem, and slide in a replacement module all without having to shut down the MBS1000 1U rackmount server.

Systems like the MBS1000 that utilize a redundant power supply by definition are modular. What is unique about the modular redundant supply approach in the MBS1000 design is that the standard auto-ranging power units can handle both 100-240VAC and 200-250VDC input voltages automatically. An option is also available for -48VDC input power exclusively. The -48VDC input power is a common voltage used in large telecom systems.

Both the data storage drive and the optical media drives are also modular. The MBS1000 chassis design enables front access, hot swap field replacement of the storage drives. The optical media bay is also accessible from the front on the 1U rackmount server; however, power will need to be shut off to replace the optical media drive in order for the operating system to recognize the new optical drive.

Now that we have a basic understand of the modular design aspects of the MBS1000 1U rackmount server, let's take a look at why these design aspects are important.

Section 2 – 1U Rackmount Server Flexibility, Longevity & Stability Considerations

Flexibility - Hardware

From a processor board perspective, flexibility is designed into the MBS1000 1U rackmount server with the support of multiple processor options that are supported with an LGA1155 socket on the server's processor blade. This approach, rather than the alternative of soldering the processor to the blade (BGA), gives the user or OEM latitude in the selection of specific processors. The socket approach also has built-in flexibility in case the application requirements change and a different processor option is required. Of course, the trade-off being that the BGA processors tend to draw less power, and run cooler; however, these BGA advantages are frequently gained at the expense of reduced processor computational capabilities. The socket approach provides more capability; plus advances in processor core thermal efficiencies, coupled with passive heat sink design improvements have made socketed processors the solution of choice in many data center applications.

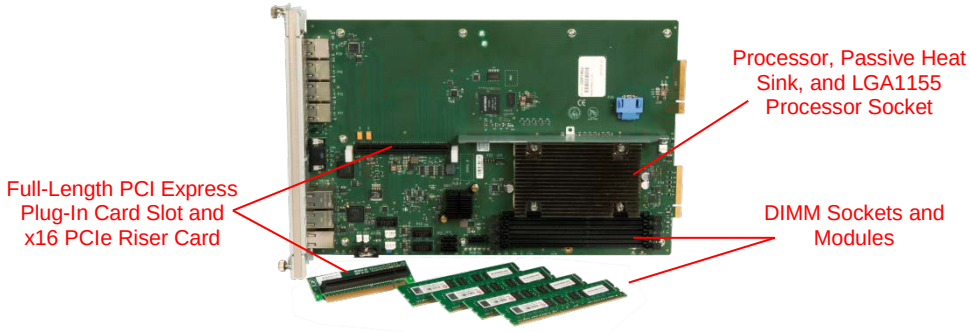

Figure 2 – MBS1000 Processor Blade

The processor blade shown in Figure 2 illustrates the approach taken with regards to system memory and PCI Express option card support. These modular design elements enable the easy addition of specific memory modules and industry standard PCI Express plug in cards in order to cover the widest data center application scenarios possible. The way the PCIe card slot is designed into the processor blade allows the end user and OEM to support a full-length PCI Express card. Many of today's PCIe cards are half-length or less; however, high-end cards such plug-in GPU cards tend to use the full-length form factor.

An alternative to plug-in memory modules in many blade cards is to solder the memory chips directly to the board. This approach can be less expensive, but in the long run is more costly because it locks the design into a specific memory capacity regardless of current or future application requirements. The memory DIMM socket board design approach is more flexible, and enables system memory scalability.

Modular power supply modules as shown in Figure 3 enable flexibility both in the types of incoming power supported, and the ability to quickly swap a failed power supply module in the field.

Figure 3 – MBS1000 Power Supply Modules

Flexibility – Software

The MBS1000 1U rackmount server is designed to be a software agnostic hardware platform in order to maximize the selection of operating systems and application software packages. The processor blade's BIOS is a menu-driven firmware package loaded into the blade's SPI devices to enable the hardware platform to be optimized for a range of operating systems and application software implementations. For example, Table 1 illustrates a short list of operating systems enabled for use on the processor blade.

Operating System Name
Windows 7.1, 32bit Pro
Windows 7.1, 64bit Ultimate
Windows 8.1 Enterprise
Windows 2008 Server, R2, 64bit
Windows 2012 Server, R2, 64bit
Linux - Centos 7.0, 64bit
Linux RHEL 7.0, 64bit
Linux- SUSE 13.2, 64bit

Table 1 – Enabled Operating Systems

In addition to operating system and application software flexibility, the MBS1000 1U rackmount server

hardware platform supports remote system monitoring via Trenton's built-in Smart System Management™ or SSM embedded application software. This software is built upon the open remote system management software standard called the Intelligent Platform Management Interface (IPMI) and its related sub-components including:

- Intelligent Platform Management Bus (IPMB)
- IPMI Platform Management FRU Information Storage Definition
- Intelligent Chassis Management Bus (ICMB)

Trenton SSM's implementation of IPMI provides seamless and efficient remote system management and control capability from any monitoring location worldwide. A short list of Trenton SMM application software functionality includes:

- Fan speed monitoring
- Fan condition & status
- Alarm monitoring
- FRU management
- Voltage monitoring
- SBC present
- Remote messaging (i.e. call home)
- Poll for processor & memory health

Flexibility – Hardware and International Certifications

Rugged card edge fingers and secure mating socket connectors are key hardware design elements that maximize the reliability of the plug-in components used in the modular 1U server design. The card edge finger approach eliminates expensive pin and socket connectors. Card edge fingers do not suffer from the alignment and pin damage associated with pin and socket connector designs. Pin and socket connectors have the advantage of greater interface connector density, and previously offered application advantages in high vibration environments. The number of interface connections required between the processor board and the server mid-plane do not require a high-density pin and socket connector approach in this modular 1U rackmount server design. Card edge contact plating improvements, and edge card secure socket connector designs work together to bring card edge connector performance in high vibration environments up to level of pin and socket connectors, but at a fraction of the overall connector cost.

Card edge fingers are much more common in the commercial-off-the-shelf (COTS) plug-in card and power supply markets. This leads to more flexibility in option card and server component choices with a corresponding increase in economies of scale of available modular components. A greater availability of COTS components reduces the overall cost of components used in a modular 1U server design that employs card edge fingers and mating sockets.

Figure 4 illustrates how the EF1 and EF2 card edge fingers of the server's modular processor blade engage with the industry standard PCI Express connectors located on the server midplane.

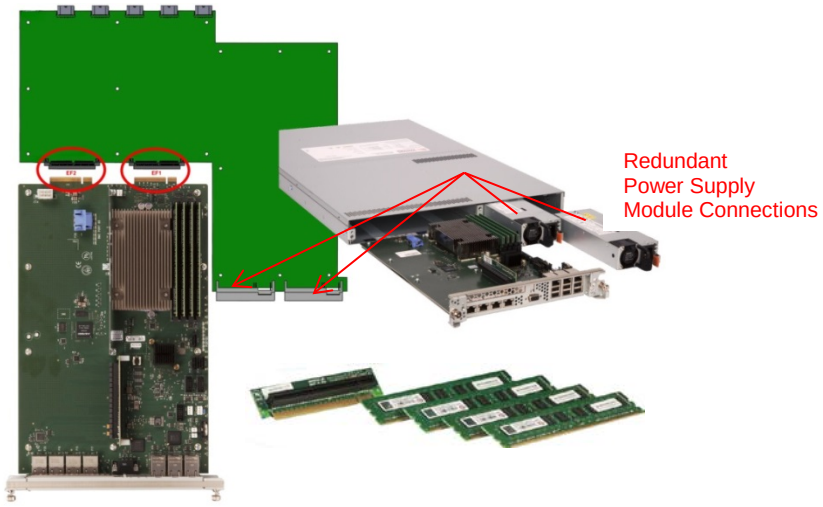

Figure 4 – Card Edge Fingers and Sockets

Similar card edge fingers and mating sockets are used for the power supply modules, PCIe riser card, and memory modules. The x16 PCIe riser card enables support for a wide variety of industry standard, PCI Express plug-in cards, while memory DIMM sockets maximize system memory flexibility.

The hardware flexibility of the modular 1U rackmount server design lends itself well to hardware platform approval by many different international certification organizations. The platform design enables easy system configuration changes in order to meet specific regulatory requirements. This aspect of the modular server design simplifies worldwide hardware deployments. Figure 5 is an illustration of the certification label for the MBS1000 with a number of international certification marks located in the lower right corner of the label.

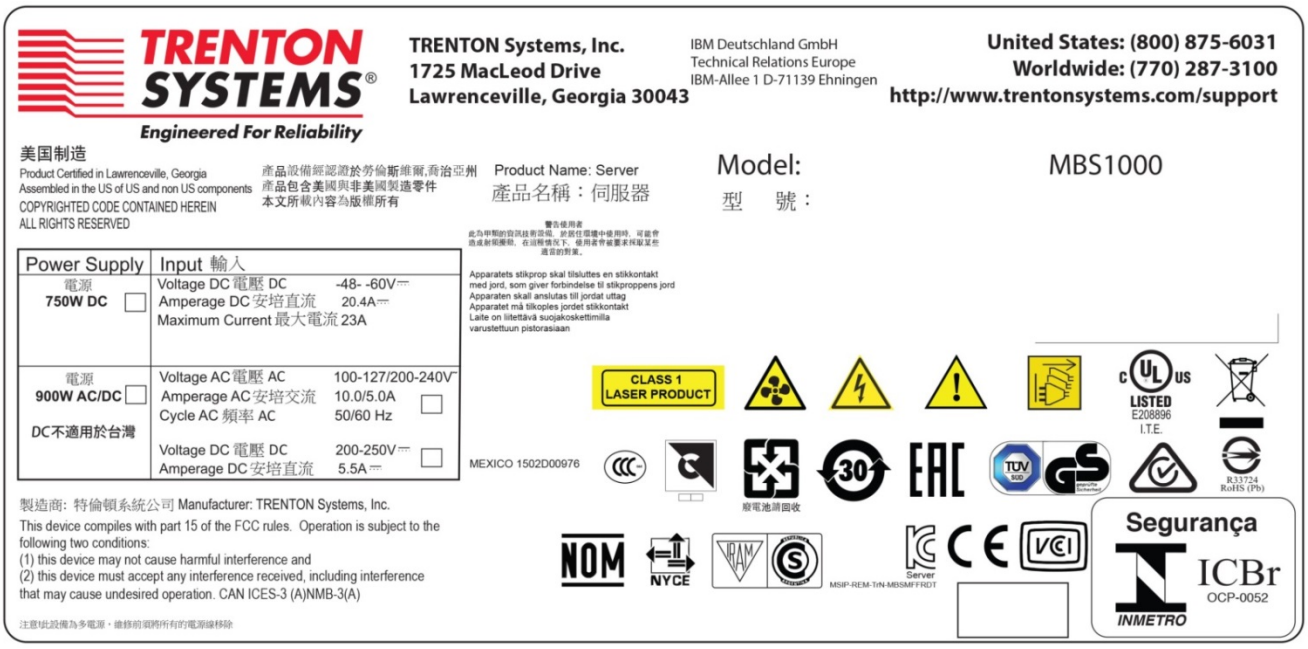

Figure 5 – Certification Label

Hardware Platform Longevity and Application Examples

One of the knocks on traditional 1U rackmount server platforms; sometimes referred to as “pizza boxes,” is the short lifecycle of the commercial processors and motherboards typically used in these platforms. The argument goes that if the pizza box goes EOL during a data center project, the 1U rackmount server is so inexpensive that you can simply throw away the old hardware and replace it with an “equivalent” platform. The replacement hardware cost may be low, but the hidden costs of data center downtime associated with changing device drivers, and making application and/or operating system software changes can be prohibitive.

The processor board used in the MBS1000 overcomes the short lifecycle issues associated with commercial processors, chipsets, Ethernet controllers, etc. by incorporating long-life i.e. “embedded” components from the Intel® IoT Group and the Intel® Datacenter Group. For example, the Intel® Xeon® processor used on the server’s processor board has a standard lifecycle availability of seven years. This extended lifecycle extends to other components on the blade resulting in a stable hardware environment for software developers. This approach eliminates the cost associated with hardware EOL-driven software upgrades.

Data Center Deployment Example – IBM z13

IBM Systems has chosen two 1U modular rackmount servers as components of their z Systems servers, the IBM z13 and IBM z13s as shown in Figure 6.

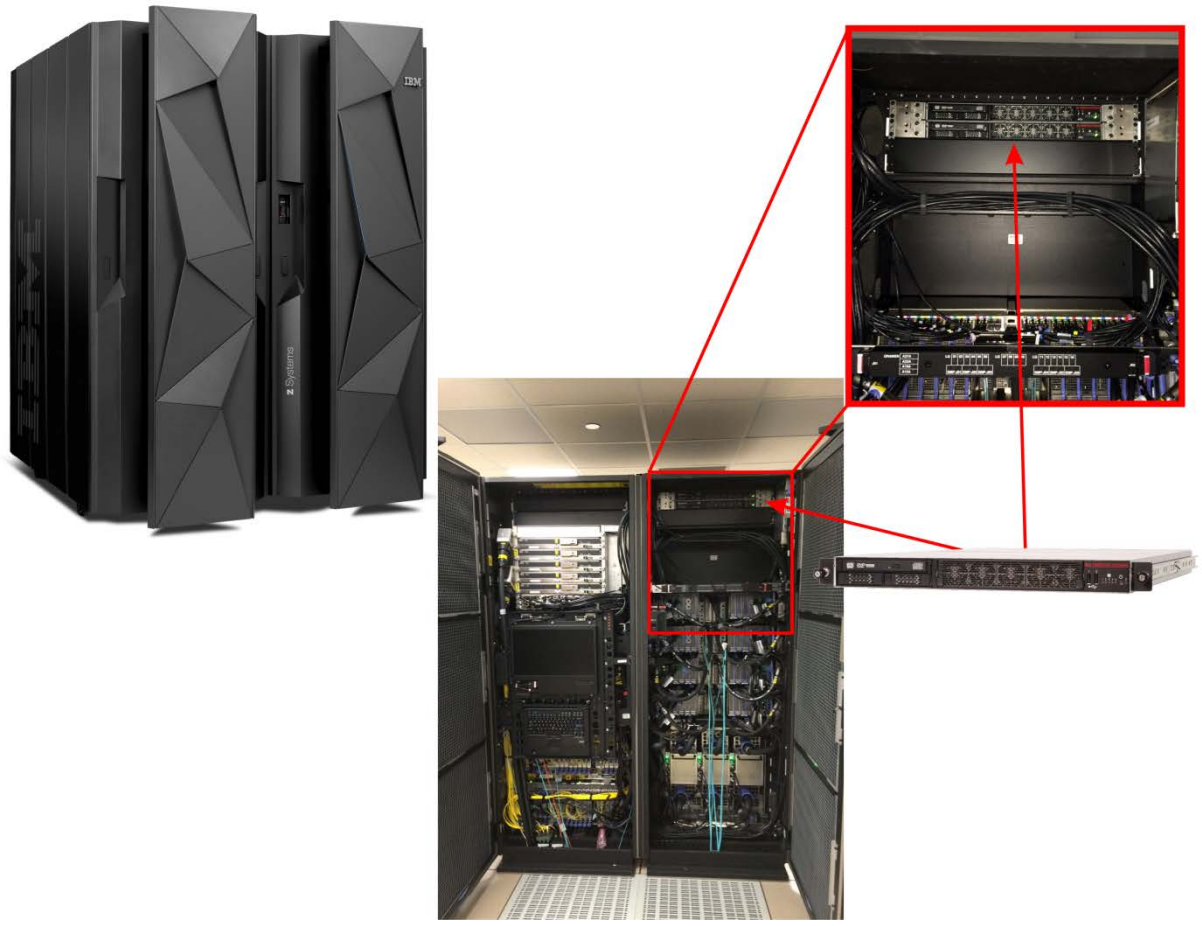

Figure 6 – Data Center Deployment Example

The IBM z13 server family offers hybrid cloud and mobile capabilities to meet the expanding data processing needs being driven by the expansion of the Internet of Things application environment. The system is designed for secure, reliable operations for reduced business risk with stronger and faster protection of critical data across a hybrid cloud environment.

The two modular MBS1000 1U rackmount servers in this data center application are not available for client use; however, the servers are key elements in the power up sequence of the IBM z13 and IBM z13s. In addition to controlling the power-up sequence, the 1U servers monitor the operational conditions of various system hardware elements, and make this data available via the dedicated management Ethernet LAN (Port 0) and the server's Smart System Management application software. System status information is also stored locally for redundancy purposes. The two 1U rackmount servers used in this application are completely redundant. The serviceability sub-section in section three provides more detail on how this redundant server configuration meets the RAS (i.e. Reliability, Availability & Serviceability) requirements of z Systems.

Other Modular 1U Rackmount Server Deployment Examples

The common application thread in all of these modular 1U rackmount server deployment examples is the need for a rugged hardware platform that is built to last. The following deployment examples are basically containerized server deployments in hostile data center environments.

- Airborne and shipboard applications for data gathering and remote communications
- Specialty vehicles use in energy exploration and production
- Forward operating bases for military operations
- Earth-bound geophysical research stations
- Earth and space-based research stations

Hardware Platform Stability

In the world of commercial computing, servers coming off the assembly line will typically experience a configuration change every six months with components swapped out and the system BIOS updated as a result. These systems also experience service lifetimes of just two to three years before being completely replaced. The realities of designing and manufacturing industrial and military-grade server solutions are quite different, as platform stability requirements demand a longer life cycle, with complete control of component selection and configuration management.

The issue of crafting long-life server solutions is two-fold; designing a product that will last for an extended period of time, and then selecting components which have extended production time frames. This approach enables a modular server configuration that can be manufactured far longer than the typical commercial server product.

As mentioned earlier, selecting long-life components such as the embedded Intel Xeon processor and related chipset that have availability horizons of 5-7 years enhances operating system and application software value. This enhanced value is achieved by reducing, and oftentimes eliminating expensive software upgrades that in the past have been triggered by a component going EOL.

Likewise, a stable 1U rackmount server BIOS that is under strict revision control eliminates costly operating system and application software deployment surprises for data center deployments that happen in the later stages of a project lifecycle. The system BIOS performs a wide range of critical functions, from setting the sequence in which devices are initialized to issuing other configuration commands. While the 1U rackmount server's processor module is designed around a standard version of system BIOS, the addition of I/O cards, such as graphics sub-systems, GPU computing cards, network interfaces or analog & digital data acquisition cards, may require a modification to the system BIOS. Having the server's BIOS under strict revision control enables Trenton Systems to provide documented BIOS modifications for a customer's specific data center implementation of the MBS1000 modular 1U rackmount server. BIOS under a documented revision control process also prevent surprises during future data center expansion by allowing the documented BIOS replication of the proven hardware and software platforms.

Section 3 – 1U Rackmount Server Reliability, Availability & Serviceability (RAS) Considerations

Hardware Platform Upgradeability Considerations

The modular design approach of the processor board is more cost efficient over time by allowing for a slide-in processor board update with future processor architecture. For example, a data center application may require different processor architecture to provide more processing cores for a specific application scenario. With the modular 1U rackmount server approach, it would be possible to design a new processor module that could slide into the current 1U chassis. Of course, there would BIOS changes to manage, but these would go hand-in-hand with the specific application software differences inherent in a new data center project. The bottom line is that there should be a need to throw out the entire 1U rackmount chassis hardware in order to support a processor architecture update.

Hardware Platform Serviceability Considerations

Extending the usable life of a 1U rackmount server enhances data center ROI in a number of different ways as discussed previously. Another way is the easy field support of the modular server. All critical components of the MBS1000 server platform can be accessed and replaced from either the front or back of the server without having to remove any chassis covers. Figure 7 illustrates the modular components on the front and back of the server.

Figure 7 – Front and Rear Field Replacement Units (FRUs) or Modules

Even the optical media drive is a front removal component via an access door located just above the HDD carriers. Table 1 lists of all of the FRUs incorporated into the modular 1U rackmount server design:

Description	Trenton Systems P/N
1U Rack Mount System w/o PS	262500003084-00
1U Base System	198500001101-00
MBC8240 SBC with Intel Xeon E3-1225v3, 32GB, Battery	198500001100-00
MBC8240 SBC with Intel Xeon E3-1225v3, w/o DIMMs, Battery	198500001110-00
8GB DDR3-1600 ECC Memory DIMM	198500001112-00
MBC8240 SBC Battery	198500001114-00
900W 100-240VAC/200-250VDC System Power Supply	198500001098-00
750W -48VDC System Power Supply	198500001117-00
Fan Module	198500001102-00
DVD with Carrier	198500001104-00
2.5” HDD Carrier	198500001105-00
1TB 2.5” SATAIII HDD with Carrier	198500001106-00
Bezel, Fan Cover	198500001123-00
26” System Slide Rail Kit	198500001127-00

Table 1– MBS1000 Field Replacement Units (FRUs)

The 1U rackmount server's modular design is ideal in meeting Reliability, Availability & Serviceability or RAS data center requirements. Virtually everything in the server design is redundant. For example, in the IBM z13 and IBM z13s deployment examples, each 1U rackmount server receives incoming power from independent power domains. The servers are cross wired so each one is aware of the other and either one can assert control.

High availability requirements are designed in by virtue of the entire communication network being fully redundant. This allows a server to tolerate a reliability failure anywhere in the network. There are also redundant signal wires to support service calls that are originated by the operating system. The 1U rackmount servers also have concurrent patch and mirroring of each server's hard drive in order to monitor the hard drive health of each server. This entire redundancy approach enables concurrent service that allows the servicing of one rackmount server while the other takes over operational tasks. One of the servers is designated as the primary and the other is the hot back-up alternate.

A mainframe service package called an MCL (Microcode Change Level), is used in the IBM z13 and IBM z13s to enable updates to the 1U rackmount server's BIOS, BMC, and other settings to meet the requirements of new mainframe usage cases. The advanced BMC utilized in the server's modular processor card also allows qualified service personnel to update the server's BIOS and BMC settings from the operating systems level.

The IPMI implementation employed in the 1U rackmount servers makes all the sensors visible to the operating system, which enables the mainframe (i.e. IBM z13 and IBM z13s, including the 1U rackmount servers) to call home and report exactly what parts may be failing. This proactive hardware analysis approach allows a service technician to come equipped with the necessary parts in hand to service the mainframe in a one trip service call. Individual FRU part numbers; such as the ones listed in Table 1 for the 1U servers, are part of the data packet included phone home network transmission. This is how the technician can order parts from the parts depot and then carry them to the client.

Section 4 – Summary and Conclusions

Utilizing a modular 1U rackmount computer design and incorporating long-life embedded components maximizes server longevity and by default software architecture stability. Flexibility is built into the modular server design both from a hardware standpoint and a software standpoint. The open platform modular hardware design is software agnostic and flexible enough to support a wide range of operating systems and application software packages. Remote server monitoring is accomplished via the server's IPMI implementation called Trenton's Smart Systems Management software. Worldwide server deployments are made possible by the various international certifications.

Hardware platform longevity offers significant ROI advantages in a wide variety of data center applications. The rugged chassis design attributes of the 1U modular server are particular useful in containerized deployments in hostile data center environments frequently found in military, energy, and field research data gathering applications. Hardware stability inherent in the 1U modular server design goes hand-in-hand with a system BIOS approach that is under strict revision control. This approach prevents deployment surprises during future data center expansions by allowing replication of the proven hardware and software platforms. The modular design of the 1U rackmount server supports system upgradability while simplifying field support.

About Trenton Systems & For Additional Information

Contact us for more information about Trenton's complete system offerings and any of our other embedded computing products or integrated computer systems.

Trenton Systems is a designer and manufacturer of military and industrial computing systems, rackmount computers, video display wall controllers, GPU computing systems, custom systems, embedded motherboards, single board computers, system host boards, and backplanes for critical embedded computing applications.

Trenton industry partnerships with Intel and other leading technology companies play an important role in the design, manufacture, integration and deployment of our high-performance system solutions. For example, Trenton Systems is an Affiliate member of the Intel® Internet of Things Solutions Alliance**.** From modular components to market-ready systems, Intel and the 400+ global member companies of the Intel® Internet of Things Solutions Alliance provide scalable, interoperable solutions that accelerate deployment of intelligent devices and end-to-end analytics.

Trenton products are designed and built in the U.S. and are well suited to deliver long-life performance, precision and reliability in a variety of mission critical solutions, from Government and Defense to Industrial Automation, Virtualization, Video Processing, Medical, Communications, Energy, GPU Computing, Test & Measurement and Video Display Walls.

Phone: 770.287.3100
Toll Free: 800.875.6031 (U.S.)
Fax: 770.287.3150
E-mail: Info@TrentonSystems.com
Web: www.trentonsystems.com

Follow Trenton on: