



Intel® Xeon® Processor Quality and Reliability Leadership

Investing in Tools and Methodologies to Improve the Data Center Experience

Table of Contents

Introduction.....	1
Quality by Design.....	2
Design for Quality	2
Early and Expanded Testing.....	2
Logic Validation.....	2
Platform Validation.....	3
At-Scale Cluster Validation.....	3
Manufacturing with Confidence	3
Manufacturing Test Flow.....	3
Screening Silent Data Errors.....	4
Customer Deployment and Support.....	5
Quality Tools for Customers.....	5
High-Quality Firmware.....	6
Conclusion.....	6

Introduction

The world’s data is growing at an exponential rate, fueled by the rapid expansion of connected devices and cloud computing. As data continues to increase in scale, greater demands on quality, reliability, and availability are placed on products and on data center infrastructure. Note that here, quality refers to correct operation under predefined circumstances. Reliability refers to the ability of a system to perform its required functions under stated conditions for a specified period. Availability refers to continuous operation to specifications without interruption.

Intel is making significant investments in tools and methodologies that deliver a high-quality Intel® Xeon® processor experience. This paper provides an overview of the Intel Xeon processor life cycle and how, during each stage of development, Intel achieves best-in-class quality. These stages include the following, as visualized in the following graphic.

- Quality by Design
- Early and Expanded Testing
- Manufacturing with Confidence
- Customer Deployment and Support



Quality by Design

Design for Quality

Intel prioritizes quality measures at each step of the design process. This process starts with a systematic processor design methodology and disciplined product definition. Intel Xeon processor quality capabilities are co-architected within the CPU silicon architecture, microarchitecture, firmware, and system software stack. This includes prioritizing [Reliability, Availability, and Serviceability](#) (RAS) features and methodologies during the design process.

Server quality is broader than simply the construction of the CPU. The entire platform interacts to impact the continuity of data center services, ultimately impacting the resulting customer experience. For example, servers rely heavily on Dynamic Random-Access Memory (DRAM) as the primary memory source for speed and cost efficiency. DRAM failures can lead to computational errors that then may go unnoticed until a server crashes.

To improve quality and reduce potential memory-error-caused server crashes, Intel implemented certain enhancements, starting with 3rd Gen Intel Xeon processors. Intel Xeon processors also provide unique RAS features to help safeguard data, automatically finding and fixing soft memory errors by detecting and correcting output data errors. More details are included in the solution brief [Improving Server Uptime in the Data Center](#).

Early and Expanded Testing

Logic Validation

Pre-silicon validation is a series of engineering processes which determine whether the product being developed meets the required specification. Additionally, pre-silicon simulation and emulation provide an opportunity to run software and firmware before physical silicon exists. This begins the process of quality optimization early in the product life cycle. During the pre-silicon stage, Intel Xeon processors are simulated with a robust test suite designed to find and fix issues that might otherwise manifest in silicon. A variety of environments and configurations are used to enhance platform stability.

Intel has a customized set of machine learning techniques—used during both the pre- and post-silicon stages—to optimize the limited run times on models and guide content to regions of the design that are difficult to evaluate. These techniques utilize both hardware and software feedback to guide testing. Silicon testing during the development phase has significantly reduced error rates throughout the Intel Xeon processor life cycle.

Platform Validation

Once integrated into a system platform, it is important that components are validated to ensure that the final platform delivers functionality and performance that will meet user expectations. Today's data center platform is a complex set of hardware and software components integrated together to meet a variety of needs. Uses range from data storage to complex Artificial Intelligence (AI) algorithms, requiring the highest compute performance coupled with high-bandwidth memory to conduct high frequency computations.

Validating platforms in an integrated fashion requires aligning mechanical, thermal, electrical, and software domains. Since hundreds of components make up a given platform, each ingredient—regardless of the cost or complexity of the component—must be optimized to ensure that the final platform works efficiently. Examples of such ingredients vary from an inexpensive capacitor to an expensive memory module. The task of platform validation is to ensure that all system ingredients interoperate correctly, providing the functionality originally intended.

Platform validation involves three focus areas:

1. Interoperability: Ensuring that platform components work together seamlessly.
2. Workloads: Testing with representative customer content or workloads.
3. Environmental: Testing with realistic environmental conditions for the platform.

When executed successfully, these focus areas help ensure that the platform delivers the performance, functionality, reliability, and experience that customers expect from an Intel Xeon processor-based system.

At-Scale Cluster Validation

Component and platform validation are augmented with at-scale cluster validation. This is meant to simulate a realistic customer environment. Intel validates different customer usage scenarios using representative workloads. In addition, Intel develops and uses fleet services similar to what a cloud service provider offers, including upgrade, maintenance, orchestration, telemetry, and data analytics. This helps ensure smooth customer adoption of Intel Xeon processors. At-scale cluster validation also serves to detect high Mean Time Between Failure (MTBF) and marginality-related failures. These failures are difficult to detect and may escape traditional validation methodologies.

Manufacturing with Confidence

Manufacturing Test Flow

Intel uses a comprehensive set of techniques at different stages of the manufacturing process to deliver products that meet stringent quality and reliability goals. These stages include the following:

1. Wafer Sort
2. Class Test
3. System-Based Test

Wafer Sort: Thorough testing of every device is first performed directly on the silicon prior to individual die being assembled into packages. This process uses a combination of industry leading embedded memory test techniques, scan-based testing, functional testing, and parametric testing to identify and discard defective die. By supplementing scan test vectors with functional tests, subtle defects that cannot be detected with conventional methods are screened, resulting in significantly improved product quality.

Intel develops functional tests for all phases of manufacturing based on both pre-silicon metrics and post-silicon analysis. Advanced analytics and machine learning methods are used to identify die that may not be sufficiently reliable, due to latent or marginal defects. Data collected about passing and failing die is fed back to Intel's wafer factories to improve the silicon fabrication process itself.

Class Test: Following wafer sort, individual die is assembled into packages. After assembly, the packaged devices receive additional stress at elevated voltage and temperature conditions so that any devices likely to fail early in the product's life can be discarded. Next, devices undergo classification testing, or class test, during which the full set of memory, scan, functional, and I/O tests are applied at the final use conditions. Class test determines the frequency and power levels at which each device can operate.

System-Based Test: The final major step of manufacturing is for every device to undergo a System-Based Test (SBT). During SBT, multiple operating systems and applications are run to verify that no defective parts escaped earlier testing. SBT hardware is based on a reference board design and includes memory Dual In-line Memory Modules (DIMMs) and a set of I/O devices. In addition to the standard workloads used during SBT, specific tests are included that screen for [Silent Data Errors](#) (SDEs).

The SDE tests used by SBT confirm that calculations performed by each device are correct. Many of these tests are part of the unique [Intel® Data Center Diagnostic Tool](#) suite, which is discussed below. Included are tests that run identical complex operations (for example, matrix math) on all processor cores in parallel, comparing the results at the end. Other tests perform reversible operations, such as encryption/decryption or compression/decompression, checking that results match the original results. Most SDE tests use pseudo-random data and instructions to maximize the ability to detect subtle, random defects that only manifest as SDE.

Screening Silent Data Errors

Failures can stem from a wide variety of sources, including radiation, aging, latent defects, logic bugs, and time-0 manufacturing circuit marginalities. These failures can then manifest as either unexpected work stoppages or, in rare cases, a SDE.

SDEs caused by manufacturing defects are typically difficult to find. The large-scale nature of data center infrastructure is such that SDEs may only manifest under specific combinations of voltage, frequency, and temperature, in conjunction with a specific sequence of operations. It is therefore essential that test methods to screen for SDEs are designed with this complex behavior in mind.

As discussed in the prior section, Intel uses many tests during SBT to screen for defects that manifest as SDE in manufacturing. Intel has demonstrated that many of the defects that manifest as SDE are not detectable through conventional Design-For-Test (DFT) methods, including scan-based and array Built-In Self-Test (BIST) tests that are used at wafer sort and class test. Targeted functional content, like those included in the Intel Data Center Diagnostics Tool (discussed below), is required to screen this important subset of defects.

Customer Deployment and Support

After Intel Xeon processors are manufactured and shipped to customers, the focus shifts to ensure high quality during the customer deployment and product life support phase. To achieve this, Intel provides both fleet management tools and regular firmware updates.

Quality Tools for Customers

Intel offers, under license, a consolidated set of tools that test for processor errors and correct functionality. Intel uses these tests in its own manufacturing processes and also makes them available for customers on the [Intel website](#). Customers use these tools in the validation of new designs, during high-volume manufacturing, and for screening in data centers. Intel collaborates with customers to understand their testing needs, updating these tools regularly to optimize their effectiveness.

The following quality tools are available for customers to use for fleet management:

- **DCDiag**, the [Intel Data Center Diagnostic Tool](#), is an exclusive tool offered only by Intel. DCDiag is designed to allow customers to test Intel Xeon processor functionality across their fleet and identify potential defects that may cause SDEs. The tool can be run as part of a regular system maintenance program, providing an easy-to-understand pass/fail result for the processors tested. This enables customers to identify potentially faulty processors during their lifetime and quickly replace them.
- The [Intel Open Data Center Diagnostics Project](#), or **Open DCDiag**, is designed to encourage industry test development collaboration. Intel recognizes there are many organizations across the industry researching ways to identify processor errors in a more effective and efficient manner. Intel founded Open DCDiag as a consistent test development framework to invite the creativity of the open-source community to enhance cloud fleet management by developing unique test screens and other innovative solutions. The project is an example of how Intel leads the industry, working continuously to improve Intel Xeon platform quality and reliability.
- To minimize server downtime, Intel provides a unique set of tools that enable customers to debug at scale and determine the root cause of issues accurately and quickly. The tools facilitate an efficient diagnosis, reducing the time to debug and enabling a faster mitigation as quickly as possible. These tools include **Autonomous Crash Dump (ACD)** and **Intel® Crash Log Technology**. Customers using a Board Management Controller (BMC) on their designs can collect debug state by enabling ACD in their fleet. Customers without BMC can rely on Intel Crash Log Technology to collect the debug state at the time of a failure. The debug state captured by these technologies can be decoded and processed by the **Crash Dump Summarizer** tool, which highlights a clear signature of the failure. In most cases, clear actions are identified that customers can take to address the failures.
- Intel developed unique **In-Field Scan** capabilities for in-fleet system screening, first available with 4th Gen Intel Xeon Scalable processors. This is like the screening that Intel does during manufacturing and helps system administrators minimize interruptions to customer operations by finding latent processor defects. Built-in test capabilities offered by In-Field Scan help detect defective cores in the field at runtime with low overhead and without taking processors offline for testing.
- Intel provides **Advanced Memory Test (AMT)** to customers to help them improve memory reliability. AMT allows customers to test DRAM health and find potentially problematic issues before impacting

end users. The AMT tool is based on test algorithms developed through partnerships with DRAM manufacturers to identify potential memory errors to improve customer manufacturing.

High-Quality Firmware

Firmware is a type of software that is embedded directly in a piece of hardware to make it work as intended. It acts as a bridge between the hardware and software running on the platform. High quality up-to-date firmware is critical for platform reliability and security. Intel releases regular firmware updates via its [Intel Platform Update \(IPU\)](#) process. Updating firmware throughout a platform's life cycle helps customers with long-term fleet management.

Due to the integrated nature of hardware, firmware, and software, product updates may require additional validation and integration from Intel's ecosystem of partners. The IPU process facilitates ecosystem coordination, leading to the release of validated updates by Intel ecosystem partners. These partners include operating systems vendors, cloud service providers, independent firmware vendors, original equipment manufacturers, and systems integrators—each of whom release validated updates to customers.

Intel addresses potential challenges related to firmware updates with the **Intel Seamless Update**, which reduces the number of system reboots required for platform firmware updates. It employs system management mode, unified extensible firmware interface runtime services, and advanced configuration and power interface services to do so. Intel also strives to make as many mitigations as possible via a microcode update which can be loaded at system runtime without reboots required.

Conclusion

Quality and reliability continue to be important considerations for Intel's data center customers. This will be more true than ever as global data and computing needs continue to grow. This paper highlighted Intel's focus on quality throughout the life cycle stages of:

- Quality by Design: Prioritizing RAS features.
- Early and Expanded Testing: Extensive logic, platform, and at-scale cluster validation.
- Manufacturing with Confidence: Screening defects through water sort, class test, and system-based test.
- Customer Deployment and Support: DCDiag, unique fleet management tools, and the IPU process.

Intel's position as an integrated device manufacturer provides a broad view across the industry, allowing greater understanding and an ability to anticipate future quality challenges. Quality is optimized end-to-end, with steps taken at each stage of the Intel Xeon processor life cycle to maximize both quality and reliability. Intel looks forward to continuous collaboration with customers to deliver on the most demanding data center quality needs.