the green grid™

ENERGY STAR® for Computer Servers Tier 2 Preliminary Draft Feedback and Recommendations (October 2009)

The Green Grid Association, a consortium of industry-leading companies, welcomes the opportunity to comment on an early draft of topics under consideration for the ENERGY STAR for Computer Servers specification.

Introduction

A consortium of information technology providers, consumers, and other stakeholders, The Green Grid seeks to improve the energy efficiency of data centers around the globe. The association takes a holistic and comprehensive approach to data center efficiency and understands that developing The ENERGY STAR® Tier 2 performance/power metric represents a significant challenge, one which requires cooperation among a wide range of industry principals. Participants in The Green Grid include such diverse companies as major server and storage equipment manufacturers, major software providers, and large data center end users/owners.

Summary

The Green Grid appreciates the EPA's efforts to capture and address the issues raised up to and including the workshop held Sept 25, 2009. With the aggressive schedule for Tier 2 for the server specification, we recommend that the focus be on development and incorporation of SPEC's Server Efficiency Rating Tool ™, limiting the scope to 1S-4S general purpose pedestal, rack and blade servers, and addressing the issues highlighted by the system manufacturers with the Tier 1 specification. While we recognize the sense of urgency, we believe that delaying in the target to April'2011 would provide the minimum amount of time to develop, test, evaluate, and prove SPEC's Server Efficiency Rating Tool™, to collect sufficient data to set the metric criteria levels, to verify and document the final testing procedures, and to incorporate the methods into the Tier 2 specification.

The Green Grid has provided proposed clarifications, and recommendations to each of the sections highlighted in the ENERGY STAR preliminary draft document, dated 9/23/09. The comments may be similar to some of the individual responses provided to the EPA, but, represent the consensus opinion of the Green Grid participants in the process.

Detailed Comments by Section

Definitions

High Performance Computer
We agree that servers that are designed in a large scale construct targeted to maximize performance for scientific research and large scale modeling should be considered HPC systems. Many of today's HPC systems reuse conventional building blocks but will vary in design from custom boards and layout with general purpose processors, to customized chassis, I/O and memory configurations using general purpose server board configurations. There has also been an increase in the use of sophisticated power management features to maximize the performance that can be attained within the larger chassis or facility configuration. The Green Grid would like to work with the EPA to refine the definition in the next draft to reflect these in an updated definition.

Overall, we believe that HPC systems should be excluded from the Energy Star Tier 2 specification due to their diverse requirements and implementations.

Resilient Server
A Resilient Server is a server designed with extensive RAS features, including error self-correction to ensure data resiliency and accuracy. Resiliency, RAS, self–correction, and accuracy

features are integrated in the micro architecture of the CPU and chipset functions in a Resilient Server. Resilient Servers are engineered with additional and more complex components in their underlying infrastructure in support of the resiliency features, which in turn require more energy to operate. For this reason, they should be distinguished from a computer server without an equivalent level of RAS features. These additional energy requirements will affect the performance of resilient servers on a performance power benchmark, potentially impacting the ability of the servers to compete effectively against servers without these features. While the intent of the development of the Active Mode Rating Tool is to provide a tool which will normalize for these features, it is possible that resilient servers will demonstrate a unique performance power signature that is fundamentally different from servers without resilient features. In this case it would be appropriate to establish specific criteria for Resilient Servers.

Definitions of the features of a Resilient Server are provided below. A Resilient Server should have all or many of these features:

- Memory Fault Detection and System Recovery: DRAM Chip Sparing, Extended ECC, Mirrored Memory
- Machine Check Architectures – Fault Isolation and Resiliency
- End to End Bus Retry
- Hot-swap components: I/O, hard drives and AC/DC power supplies
- Ability to perform on-line expansion and retraction of hardware resources without OS reboot - also referred to as "on-demand"
- Multiple physical banks of memory and I/O adapters

Storage Equipment

The Green Grid agrees that the definition should be consistent with the definition in the ENERGY STAR Storage specification, and we fully endorse the recommendation provided by SNIA.

Processor Utilization

Though "processor utilization" was referred to in the Green Grid whitepaper, we suggest the following changes to its definition. In particular, processor utilization is a percentage figure representing how busy the processor has been relative to the rated frequency of the unit. Although there are modes of operation where the frequency may change to be temporarily higher or lower, the relative measure provides the level of activity against the planned peak operation. For monitoring purposes, the relative percentage value allows the end user or system management systems to better appropriate the computing resource to meet upcoming demand.

The definition of processor utilization should be changed to read: "The percentage estimate of the server's compute activities relative to the full operation specified voltage and frequency of the processor(s)."

EPA's Questions

1. *What additional terms and definitions should be added to the Tier 2 specification? What is the anticipated effect (if any) on existing program scope and requirements?*

None. The additional definitions and changes reflect those systems which are designed and configured explicitly for specialized activities. The specialized configurations and designs of the new proposed product descriptions have energy profiles that do not allow for an equivalent energy comparison to general purpose servers.

2. *Are changes to existing Tier 1 definitions suggested for Tier 2? What developments or factors should EPA be aware of that suggest the change? What is the anticipated effect (if any) on existing program scope and requirements?*

We have proposed changes to the specialized systems definitions to better describe them. The intent is that they be treated as a separate class of system with separate criteria rather than being compared to general purpose computers.

3. *With the goal of defining a broad taxonomy of server categories in the Computer Server Types section, what server categories are missing from the current list provided? Do the definitions define a continuum of categories present on the market? Are existing definitions clear enough to avoid overlap between categories?*

The distinctions are clear enough. Although future overlap or new system categories may not be completely avoidable, the distinction is sufficient for the current cadence of the Energy Star specification.

4. *Are there any other sources that the EPA should review for variations of, or additions to, this list of definitions?*

Other sources that may track the various definitions of computer servers, as markets and system devices expand include ATIS, SNIA, CSCI, ECMA, Digital Europe, IDC, Gartner, and others.

5. *As "Active Mode" is an anticipated area of effort for Tier 2, are there any standardized or agreed to characteristics that could be referenced as building blocks of an Active Mode definition?*

Energy consumption in "Active mode" is characterized by the system configuration, compute capacity, extensibility, power management features along with its resume latencies, and system state residencies. These describe some of the salient features that should be covered by the active mode evaluation tool, categories, and the product information disclosures in the Tier 2 specification.

Qualifying Products

Due to the existing market segmentation of servers, we believe the focus should remain on 1S-4S servers and the inclusion of bladed systems. As indicated by 2009 IDC market data, over 98% of the systems shipped are within these categories. Given the aggressive schedules we recommend focusing on the Server Efficiency Rating Tool™, incorporating blades, and addressing the issues highlighted in the Tier 1 specification.

Server Appliances, >4S Servers, Fully Fault Tolerant Servers, and Multi-node Servers offer unique features and characteristics that operate under a significantly different energy profile than the general purpose 1S-4S servers. Until such time that the evaluation method can be developed that would distinguish energy efficient variants of similar systems in these categories, we do not recommend including them in scope of the Tier 2 specification.

Expansion of the program beyond 1S-4S general purpose servers should not be considered until after the Server Efficiency Rating Tool™ and ENERGY STAR procedures are proven for the larger class of 1-4S general purpose servers. For those systems that can not be included in Tier 2, we recommend an advisory to the procurement agencies to treat these classes of machines as out of the scope for the current revision of the ENERGY STAR for Server program. Customers of those classes of machines should not use ENERGY STAR as a purchasing criterion, until such time as these categories of product can be considered in scope. (i.e. a future revision of the Energy Star for Server program).

EPA's Specific Questions:

Discussion Topics: *> 4 Socket Servers, Server Appliances, Fully Fault Tolerant Servers and Multi-Node Servers*

1. Is there new information on any of these categories (increased market share/interest, test procedure availability, comparable sets of products, market differentiation/energy savings potential) that supports further investigation by ENERGY STAR?

We do not believe there have been significant changes in either the volume of shipments, available testing procedures, or energy profile information that support prioritizing investigation for incorporation to the ENERGY STAR program at this time.

2. EPA received an initial suggestion that the Tier 2 specification continue to focus on volume servers (1S, 2S, or 4S general purpose servers in rack, pedestal, blades, or chassis form factors). With the exception of addressing blade servers, this suggested scope is similar to the initial tier of the program, and includes the majority of products on the market. Are there suggested areas of the market outside of this scope – including the four product classes noted above – that have a critical mass of products to allow effective comparison, represent a large source of energy-saving opportunity, or otherwise present an opportunity for ENERGY STAR to differentiate the market?

We do not feel that the scope has largely remained the same. The inclusion of blades and having a metric evaluation for 3S and 4S systems represent a significant scope increase when the creation of a metric to evaluate their general efficiency is taken into account.

3. What is the relevance of servers described by the new definitions for *Resilient Servers* and *High Performance Computing Systems* to the overall scope of the ENERGY STAR Computer Servers program?

For the evaluation of the general purpose systems, the new definitions are machines whose energy profile may differ depending on the capabilities of the metric used. For Resilient Servers, we recommend that these systems be reviewed to determine if they should be evaluated within the same categories as 1S-4S systems, should contain adjustment factors, or should be considered another category in the Tier 2 program. For HPC systems, we believe the energy profile and specialized function make HPC systems distinctly separate from the general purpose machines. We noted that HPC systems should be considered out of scope for Tier 2, given these characteristics.

Approach: Blades

As discussed in the 9/25/09 workshop, we believe there are evaluation methods that can incorporate Blade Systems in the program for Tier 2. Extending development of the Tier 2 requirements and evaluation methods to include Blade servers is preferable to attempting to develop an interim evaluation method. Developing a new evaluation method as the interim Tier 1 process, as suggested in the preliminary draft specification is unlikely to offer any coverage given the Tier 2 schedules, provides little reuse for Tier 2, and may delay Tier 2 development.

If an interim Tier 1 qualification requirement is needed, we recommend treating blades similarly to 3S and 4S systems. Treating blades in Tier 1 with similar requirements as 3S and 4S systems allows inclusion of blades systems in the Tier 1 specification without an extensive data creation, collection and evaluation effort to set a Tier 1 idle power criterion, would promote blades as a consolidation solution as data centers refresh their equipment and provide power performance information that could be used by customers prior to the implementation of the Tier 2 specification.

Tier 2 Proposal for the Blade Servers:
We propose that once a blade chassis and server combination are qualified under ENERGY STAR, a qualified product would consist of a chassis and blade servers from qualified product families. This is because customers will mix and match blade servers in a chassis based on their computing needs.

This approach offers the most effective means to bring blade servers under the ENERGY STAR computer server requirements given the complexity of the blade systems and the varied approaches to distributing cooling and power distribution overhead between blade servers and blade chassis by different manufacturers. This approach has the following benefits:

1. By testing a full chassis (or a partially populated chassis if the manufacturer prefers) of minimally configured blades, data will be provided to allow calculation of per blade power usage in a fully populated chassis. The minimum blade server configuration was chosen for this test, as it provides the best means to allow meaningful comparisons between manufacturer's systems. It is likely that typical and maximum configurations will vary sufficiently in power use to make it difficult to derive meaningful comparisons between systems.
2. It will be possible to estimate the chassis power by subtracting the power use from a minimally configured blade server (times the number of blades in the chassis) tested in a test fixture or in a single chassis slot from the power measured by the fully populated blade system.
3. It highlights the range of the power levels from a minimum to maximum blade server configuration as tested in a test fixture. As such, a data center manager would be able to determine range of power requirements needed for the product family and its applicability within their power constraints.

It is our belief that testing the individual servers and the full chassis provides the data necessary to allow a data center operator to easily calculate and compare the power use of different configurations and levels of chassis population. If a manufacturer chooses to test a partially populated blade chassis to compute the per blade server power use, that should be allowed as the per blade power use should be conservative due to distribution of the chassis overhead over less than the full number of blades. The proposed approach provides a workable testing regime for the manufacturers, providing the data needed by the server purchaser to make informed decisions on the blade system power use.

Added Definitions for blades:

Maximum Configuration; Blade Chassis: A Blade Chassis populated with blades in the maximum configuration for that product or product family.

Minimum Configuration; Blade Chassis: A Blade Chassis populated with blades in the minimum configuration for that product family.

Typical Configuration; Blade Chassis: A Blade Chassis populated with blades at the typical configuration for that product or product family.

Multi-wide or Tall Blade Server: For any blade which is larger than one bay, the size represented by that unit will be equal to the number of bays occupied. If there are multi-wide (or tall) blades in a chassis, the capacity of the chassis is reduce by the number of bays occupied by the multi-wide/tall blade minus one. (i.e. a 3 wide blade will reduce the chassis capacity by 2 slots).

EPA's Blade Discussion Questions:

1. *What efficiency/power/performance information do purchasers routinely request when investigating a new blade system? Do these requests change at all if the purchase is intended to replace standalone servers rather than other blades?*

Manufacturers provide data showing that power per server is lower for comparable blade servers relative to rack-mount servers. At this time, there are no standard benchmarks that are used by customers to compare rack-mount servers to blade servers.

2. *What efficiency/power/performance information would be useful to blade purchasers that is not routinely requested that could influence the provisioning process?*

Blade customers tend to know what to ask.

3. *What assumptions must typically be reported when marketing comparative blade performance and efficiency?*

Using a CPU burn application like Prime95, companies have been marketing the relative power consumption of blade server solutions versus rack mount server "equivalents". This particular test is not a fair comparison, however, since there is no performance aspect in the test. There is some hope that SPECpower and other benchmarks will soon provide a fair means to compare blades and rack mount servers.

4. *Regarding infrastructure overhead (power distribution/supply, cooling), how can the most efficient implementations be identified? What assumptions are fair/relevant to ensure fair comparison?*
    a. It is preferable that an enclosure <u>not</u> have its own ENERGY STAR compliance process. We have suggested a means for qualifying blades in certain enclosures.
    b. An easy and low cost approach to measuring the power of the enclosure/chassis infrastructure would be to measure the power with one server in an enclosure and then measure power with two servers installed, then subtract twice the power difference to negate the server blade power A benchmark should run on both servers during the power measurements, so that there is nothing hibernating. The value to the customer of this measurement is questionable, however, as it introduces complexity to the rating system that may exceed the value of the metric.

5. *Given the server focus of this specification, what are ways that blade storage and network equipment could be addressed to create stable testing conditions between competing implementations?*
    a. Compliance testing of blade servers should not require the presence of storage and networking blades.
    b. The presence of storage and networking blades in an as-shipped solution must not affect a blade server's Energy Star compliance.
    c. Create adders for networking or storage blade components or exclude the components entirely for the Energy Star systems compliance requirements.

6. *Is analysis at the chassis level a valid approach to determining requirements for blades?*
    Testing the aggregated power and performance of a chassis with blades in it is the appropriate way to test compliance. However, the compliance test set-up would be representative of the deployed systems and not necessarily the exact as-shipped configuration. A testing approach has been proposed on pages 5 and 6.

7. *Are there any anticipated purchasing practices when a user moves to a blade architecture (e.g. customers typically purchase blade chassis fully populated, half populated, etc.)?*
    It depends. If a customer buys one enclosure, then it is probably partially-populated, but if they buy many enclosures then many (but not all) are probably fully-populated. Partially-populated enclosures are more likely over time to be heterogeneously populated with dissimilar blades.
    An enclosure is considered an option that is purchased with a blade server, though a specific model of a blade server will only operate within a specified set of enclosures. This doesn't preclude a blade from being bought without an enclosure option and plugged into an existing enclosure.

Active Mode Efficiency Rating Tool

The Green Grid and its members fully support the current effort by SPEC Power to generate the Server Efficiency Rating Tool ™. This tool currently offers the best approach to establishing an active mode rating for computer servers. The SPEC efficiency rating tool is an industry initiative with broad support and should provide a robust efficiency rating tool which incorporates both active and idle mode energy characteristics of a server. Given this comprehensive energy

assessment, a separate idle specification is not recommended and may conflict with the rating tool's assessment. We propose support of the SPEC efficiency rating tool as a priority for the Tier 2 Energy Star for server specification.

We recommend the following modifications to rating tool's targets which should enhance the applicability of the tool and simplify the overall evaluation of energy use by computer servers:

1. The rating tool should focus on performance and power capabilities of the different key functions which are common to all servers including, but not limited to, processor computation, both floating point and integer computations and disk, I/O, and memory access. A tool which focuses on functional, energy efficiency performance offers the best approach to providing an effective comparison of systems with both lighter and heavier configurations and higher powered processors.
2. Idle power requirements should be included as part of the evaluation under this metric, allowing the server to receive a single, overall score which incorporates both active and idle power performance. Including idle in the metric provides the opportunity to normalize the overall score for more capable (higher powered) systems which are currently penalized in the Tier 1 "idle" only criteria.
3. The tool should be designed to allow manufacturers to demonstrate their hardware and system software based efficiency capabilities. Manufacturers have enabled their hypervisors, operating systems, and middleware with capabilities that enhance the energy efficiency of the server operation. These capabilities and functions should be recognized for their ability to provide the most energy efficient solution for a customer.

The Green Grid recognizes the technical challenges the SPEC working group may have in meeting EPA's timeline of a published Tier 2 specification in October of 2010. We recommend the EPA consider extending the timeline by up to six months pending SPEC demonstration of sufficient progress and technical potential during the first half of 2010. We believe this plan is advisable for several reasons, including:

1. Currently, there are a limited number of computer server products qualified to the ENERGY STAR Computer Server requirements, indicating that the current Tier 1 specification is already challenging and separating the servers currently available in the marketplace.
2. While Option D (reference: EPA Active Mode Discussion Topics, 9/25/09) may be attractive, it will result in the ENERGY STAR computer server specification going through three very different criteria sets in the space of two to three years. This creates a moving set of standards which impact design and testing requirements. Given the long development (>3yrs) and deployment cycles for servers (>3-5yrs), design changes will have not appear until well after changes prescribed in the multiple revision plans.
3. Option D (and B) also suffers from the difficulty of selecting a representative set of benchmarks with both power and performance report aspects. On the one hand, the set of suitable, released benchmarks is still small. On the other hand, the area is in flux as a number of benchmarking organizations begin to consider the importance of measuring energy-efficiency. In addition, any benchmark suite constructed under Option D will lack a uniform infrastructure and a consistent set of rule run rules. The industry will find it difficult to execute the number of tests likely to be required.
4. Existing performance benchmarks are strictly catered towards very specific, configurations optimized for that particular workload. Each performance benchmark is used to advertise the peak capability for that particular model of system configured and optimized for that function. Running a number of different workloads on fixed configurations falsely reports the capability of that family of product and leaves the additional difficulty of merging dissimilar performance units. Though there may be methods to resolve these issues, such activity would incur delays in other programs.
5. Significant resources are required to test equipment currently available in the marketplace to new or revised server criteria. The proposed active mode testing, whether done through the

rating tool developed by the SPEC Power working group, or using existing workload benchmarks, is not done on a routine basis for all product families. It will be difficult to generate benchmark performance data for one iteration of the Tier 2 specification, and nearly impossible to do it for two iterations within 12 months.

If it is determined that the SPEC Power "Server Efficiency Rating Tool" cannot meet a reasonable schedule or exhibit the required characteristic, we recommend an evaluation approach that would utilize a select number of current benchmarks to evaluate server Active Mode power. The Green Grid recommends and would fully support investigation of either a rating system per benchmark from a select set, or a method to integrate and normalize the performance scores on a fixed set of benchmarks. We believe the investigation to these options could begin early next year pending the progress made by the SPEC team.

Regarding Tier 2 requirements, The Green Grid believes that servers originally qualified under Tier 1 should remain qualified (grandfathered) for some period of time after Tier 2 takes effect, for instance for one year or until the server model is discontinued, whichever is less. This will encourage manufactures to qualify servers under Energy Star as soon as they are available, without the additional burden of having to re-qualify all existing qualified servers as soon as Tier 2 takes effect. Alternatively, we recommend that the EPA allow servers to qualify under the Tier 2 rules in the 9 month period between the Tier 2 specification being finalized and the effective date, and that these servers automatically also qualify for immediate Energy Star Tier 1 certification. This would allow the manufacturers to test and submit their servers once during this period, for consumers to get the energy efficiency information, and eliminate the requirements for double testing and submission.

Power Supply

Net Power Loss (NPL)
The Green Grid fully supports the CSCI recommendation to maintain the current efficiency method for evaluating power supplies. The Green Grid does **not** support the incorporation of the net power loss approach, as either a criterion or data sheet reporting requirement. Net Power Loss is a fair description of the loss at low to no-load conditions. As the load increases the efficiency percentage is a much more appropriate measure over the typical operating range of a server. The industry and its customers are familiar with and understand power supply efficiency and what it means to power consumption. A change in metric would require training on terms and representative values, measurement protocols, and what the metric means to the end user. Given the limited information that data center operators will glean from the NPL and the additional cost and effort to the industry supply chain, we recommend dropping the Net Power Loss approach as a testing criteria or reported value.

Significant improvements in power supply efficiency are reaching the point of diminishing returns. On server types that have performance per watt benchmarking pass/fail criteria, there is little to no value to adding power supply efficiency requirements. With the benchmark, the server can be viewed as an independent entity. The removal of power supply testing, once a performance for system watt benchmarks are in place, supports the ability for this specification to sustain future innovations to power supply and power deliver architectures.

Power Supply Efficiency
Power supplies are long lead items for servers, are usually custom designed for a product or group of products, and require extensive worldwide safety and regulatory certifications. The expense of these power supplies is such that the continuous revision of the power supply efficiency requirement in the standard is a cost burden for industry. We recommend keeping the existing power supply efficiency limits for the next Tier.

The multi output power supply efficiency should line up with the Energy Star 5.0 requirements for small scale servers.

Energy Efficient Ethernet

The requirement of implementing the Energy Efficient Ethernet standard upon release would have significant business impact. The standard is not released or considered final and silicon and software impacts are not understood. It is not reasonable to require an unfinished, future standard in Tier 2 without understanding the implications of the standard to product performance and availability. The Energy Efficient Ethernet standard should only be incorporated into the Tier 2 after the Ethernet specification has been finalized, interoperability proven and the implications of its inclusion in the ENERGY STAR criteria can be comprehended.

Real Time Reporting Requirements

Processor Utilization Measurements: The computer server will provide an estimation of the processor or system utilization that is visible to the operator or user of the computer server through the operating environment (operating system, hypervisor, or other management interface). The number is intended to provide the data center operator a qualitative indication of the amount of load on the system to provide guidance on opportunities for virtualization and consolidation of workloads. In systems that support active and inactive (powered off) processor states, it is only a measure of the active processors.

Currently, the industry does not have a single, standardized methodology for measuring processor utilization, particularly in virtualized environments where there are simultaneous threads being managed by the processor. TGG proposes that the utilization requirement continue to be qualitative in nature for the same reasons cited during the review of the same language in the Tier 1 specification as well as some additional reasons which are specific to current technology and the nature of the Tier 2 specification.

For the purpose of the Tier 2 ENERGY STAR® computer server requirements, any utilization measurement needs only be sufficiently accurate for the purpose of enabling decisions around the consolidation or re-provisioning of workloads or the movement of virtual machines. Existing CPU utilization measurement algorithms are already sufficiently accurate for this purpose. Many data centers today use CPU utilization numbers as reported by currently shipping operating systems on existing servers. They use this information to dynamically manage data center power consumption by re-provisioning workloads to minimize the underutilization of machines. If a change to CPU utilization is needed, the industry would like to understand the reasoning for a the change in approach to CPU utilization, and how this change of approach benefits the Energy Star program or data center managers in management of the data center.

Standard Information Reporting Requirements

The Power and Performance Data Sheet and QPI forms require the same information in two different formats. (i.e specifically PSU Efficiency and PFC values) Also many of the other entries are duplicated. A consolidated sheet should be considered. For family declarations the QPI improperly calculates maximum or minimum configuration values for system memory or disk capacity. Families that have common form factor but different wattage power supplies are required to identify two families, for example.

There are literally only a few data sheets available for review from OEMs on the ENERGY STAR web site. The section outlined by line 554 would seem to indicate that complaints have been registered about the documents currently available. We would like further clarification to understand is the purpose of the information being delivered. In addition the complexity of

reporting through the power and performance data sheet and the QPI is such that there are many opportunities to have misunderstanding or improperly listed products resulting from the programming in the sheets.

Conclusion

The Green Grid anticipates a successful and collaborative development of the ENERGY STAR for Computer Server Tier 2 specification with all industry stakeholders and the EPA. We believe with the focus on the volume categories of servers and a metric that tracks the maximum performance within an energy envelope, the Tier 2 specification can be very successful. The combination and consistency of the ENERGY STAR for Computer Server program and the efficiency initiatives in the EPA and US DOE should help in accelerating the efficiency in operation of the data center. The Green Grid will continue to collect industry-wide inputs to work with the EPA in developing the ENERGY STAR programs on ICT equipment. Please feel free to contact us both to clarify and collaborate on the development of the draft, the specifications, and the implementation of the program.

.