

ORNN MASTER SERVICES AGREEMENT

This Master Services Agreement (this “MSA” or “Agreement”) is entered into as of _____ (the “Effective Date”) by and between Ornn Clearing LLC, a Delaware limited liability company (“Ornn” or “Customer”), and

_____, a _____ (“Provider”). Ornn and Provider may each be referred to as a “Party” and together as the “Parties.”

This MSA establishes the general terms under which Provider may provide GPU and other compute infrastructure and related services to Customer from time to time. The specific infrastructure, capacity, pricing, and term for each engagement are set out in one or more Orders that reference this MSA. The service levels applicable to the infrastructure are set out in the Service Level Agreement and Working Node Standard attached as Schedule 1.

1. Structure of the Agreement; Orders

1.1 General Framework. This MSA sets out the recurring legal and commercial terms that govern Provider’s provision of compute infrastructure and related services (the “Services”) to Customer. This MSA does not by itself obligate either Party to purchase or provide any Services.

1.2 Orders. Each specific engagement is documented in an order form executed by both Parties that references this MSA (each, an “Order”). Each Order sets out the deal-specific terms for that engagement, including the reserved infrastructure, node and GPU counts, hardware profile, location, pricing, billing schedule, delivery date, testing period, and service term. No Order is binding until signed by both Parties.

1.3 Incorporation. Each Order incorporates, and is subject to, this MSA and the Service Level Agreement and Working Node Standard attached as Schedule 1 (the “SLA”). This MSA, the SLA, and a given Order together form a single agreement with respect to the Services described in that Order (an “Engagement”). Terms of one Order do not affect any other Order.

1.4 Order of Precedence. With respect to a given Engagement, in the event of a conflict: (a) the body of the applicable Order controls with respect to commercial terms, pricing, payment, term, capacity, delivery, and Party-specific obligations; (b) the SLA controls with respect to service levels, availability measurement, working-node requirements, downtime measurement and remedies, maintenance, monitoring, and support targets; and (c) in all other respects, this MSA controls. An Order may modify a term of this MSA or the SLA only for that Order, and only where the Order expressly identifies the section being modified.

1.5 Defined Infrastructure. The compute capacity reserved for Customer under an Order is the “Infrastructure” for that Order. References in this MSA and the SLA to the “Infrastructure” or the “Services” mean the Infrastructure and Services described in the applicable Order.

2. Definitions

Capitalized terms have the meanings given where first defined in this MSA, in an Order, or in the SLA. In addition: “Service Commencement Date” means, for an Order, the date and time the Infrastructure is delivered and made available to Customer through the agreed access method; “Service Term” means, for an Order, the period during which the Infrastructure is reserved as stated in that Order; and “Testing Period” has the meaning in Section 4.4. Terms defined in the SLA (including “Downtime,” “Working Node,” “Availability,” and “Service Credit”) have the meanings given in the SLA when used in this MSA or an Order.

3. Provision of Infrastructure

3.1 **Reserved Capacity.** Subject to Customer's payment obligations and compliance with this Agreement, Provider shall provide Customer with reserved access to the Infrastructure described in each Order. Unless an Order states otherwise, the Infrastructure is provided as dedicated compute capacity.

3.2 **Provider Responsibilities.** Provider is responsible for delivering and operating the Infrastructure in accordance with the applicable Order and the SLA. Provider is not responsible for Customer's application development, model training, data preparation, workload administration, or software licensing unless expressly stated in an Order.

3.3 **Changes to Reserved Capacity.** The reservation under each Order is based on the node count, GPU count, location, and hardware profile stated in that Order. Any change to the committed node count, GPU count, location, or material hardware profile requires the written agreement of the Parties, documented in an amended or new Order.

4. Access, Delivery, and Acceptance

4.1 **Delivery.** Provider shall use commercially reasonable efforts to deliver and make the Infrastructure available to Customer by the target delivery date stated in the applicable Order, subject to receipt of any required payment and completion of reasonable onboarding requirements.

4.2 **Customer Responsibilities.** Customer is responsible for its accounts, credentials, SSH keys, software environments, workloads, data, models, backups, and output generated through use of the Infrastructure.

4.3 **Access Method.** Customer shall access the Infrastructure only through the access method agreed by the Parties in the applicable Order or otherwise in writing, including any required SSH, VPN, bastion, portal, or other approved access path.

4.4 **Testing Period.** Notwithstanding anything to the contrary, upon Customer receiving usable access to the Infrastructure, all required credentials, and the applicable validation results required to be provided under Appendix A of the SLA, Customer shall have a testing and evaluation period (the "Testing Period") of the duration stated in the applicable Order (and, if none is stated, five (5) consecutive calendar days) to verify that the Infrastructure conforms to the specifications in the Order. If, on or before the end of the Testing Period, Customer notifies Provider in writing that the Infrastructure has materially failed Customer's testing, the affected Order shall terminate with no further obligation of either Party and Customer shall owe no fees for that Order. If Customer does not provide such notice on or before the end of the Testing Period, the Infrastructure shall be deemed accepted, the Order shall become fully effective and binding, and the Service Term shall commence on the Service Commencement Date. No service fees accrue during the Testing Period.

5. Service Levels and Support

5.1 **SLA.** The service levels, working-node requirements, maintenance rules, downtime measurement, service-credit process, support severity levels, exclusions, and related operational standards applicable to the Infrastructure are set out in the SLA (Schedule 1), which is incorporated into this Agreement and applies to each Order during its Service Term.

5.2 **Conflict.** If an Order or this MSA conflicts with the SLA, the order of precedence in Section 1.4 applies.

5.3 **Exclusive Remedy.** Service credits or other SLA remedies, if any, are Customer's sole and exclusive remedy for ordinary availability failures, except to the extent otherwise required by applicable law or expressly stated in this Agreement or an Order, and subject to Section 8.5 of the SLA.

6. Financial Terms

6.1 **Fees.** Customer shall pay the fees stated in each Order. Unless an Order states otherwise, fees are based on reserved capacity for the Service Term, calculated at the per-GPU-hour or other rate stated in the Order.

6.2 Reserved-Capacity Billing. Fees are calculated based on the committed reserved capacity for the Service Term, whether or not Customer fully uses the Infrastructure, unless Provider fails to make the Infrastructure available in accordance with the applicable Order.

6.3 Invoicing and Payment. Each Order sets out the payment terms for that Order, including any upfront deposit or prepayment, the amount and timing of invoices and installments, whether amounts are payable in advance or in arrears, and the applicable payment due date or net terms. Provider shall invoice, and Customer shall pay all undisputed invoiced amounts, in accordance with the payment terms and billing schedule stated in the applicable Order. If an Order does not specify payment terms, Provider shall invoice the service fee as per the payment terms in the order form.

6.4 Taxes. Fees are exclusive of taxes. Customer is responsible for all sales, use, value-added, and similar taxes arising from an Order, excluding taxes based on Provider's net income.

6.5 Disputed Amounts. The Parties shall work in good faith to resolve disputed amounts promptly.

7. Term and Termination

7.1 MSA Term. This MSA begins on the Effective Date and continues until terminated in accordance with this Section 7. This MSA remains in effect for so long as any Order remains in effect, even if the MSA would otherwise have terminated.

7.2 Order Term; Renewal. Each Order begins on its Service Commencement Date and continues for the Service Term stated in that Order, unless earlier terminated in accordance with this Agreement. If, and only if, an Order expressly provides for automatic renewal, that Order will renew at the end of the then-current Service Term for successive renewal terms, each equal in length to the initial Service Term unless the Order states a different renewal period, on the same terms except as the Order or this Section provides. Either Party may prevent renewal by giving the other Party written notice of non-renewal at least thirty (30) days before the end of the then-current term, or by any different deadline stated in the Order. Provider may adjust the pricing applicable to a renewal term by giving Customer written notice before the start of that non-renewal notice period; if Customer does not wish to continue at the adjusted pricing, Customer may elect not to renew under this Section. An Order that does not expressly provide for automatic renewal does not renew, and any renewal, extension, expansion, reduction, or location change of such an Order requires the written agreement of the Parties.

7.3 Termination for Material Breach. Either Party may terminate this MSA or any affected Order upon written notice if the other Party materially breaches this Agreement and fails to cure the breach within ten (10) days after receiving written notice describing the breach in reasonable detail. If a material breach is incapable of cure, the non-breaching Party may terminate upon written notice.

7.4 Termination for Insolvency. Either Party may terminate this MSA and any Order upon written notice if the other Party becomes insolvent, makes an assignment for the benefit of creditors, admits in writing its inability to pay its debts as they become due, becomes subject to any bankruptcy, receivership, liquidation, dissolution, or similar proceeding that is not dismissed within thirty (30) days, or ceases or threatens to cease carrying on business in the ordinary course.

7.5 Effect of Termination; Refund of Unused Reserved Capacity. Upon expiration or termination of an Order, Customer's right to access the corresponding Infrastructure ends and Provider may disable Customer's access. If an Order is terminated before the scheduled end of its Service Term, Provider shall refund any prepaid fees allocable to reserved capacity after the effective termination date, calculated pro rata using the reserved hourly cost stated in that Order, within thirty (30) days after the effective termination date. Provider may offset unpaid amounts then due, and Provider is not required to refund prepaid fees if Provider terminates due to Customer's uncured material breach or Customer's insolvency, except to the extent required by applicable law. Termination of one Order does not terminate this MSA or any other Order unless expressly stated.

8. Use Restrictions

8.1 Lawful Use. Customer shall use the Infrastructure only for lawful business purposes and in compliance with all applicable laws and regulations.

8.2 Prohibited Activity. Customer shall not use the Infrastructure for unauthorized access, credential theft, malware distribution, denial-of-service activity, network interference, illegal content processing, sanctions evasion, or any other prohibited or abusive activity.

8.3 Suspension. Provider may suspend access immediately if Provider reasonably believes Customer's use violates this Section 8 or creates a security, legal, or operational risk. Provider will use commercially reasonable efforts to limit any suspension to the affected use, workload, account, or access path where practicable.

9. Data and Security

9.1 Customer Ownership. Customer retains all right, title, and interest in Customer's data, models, software, configurations, and workload outputs. Provider does not acquire ownership of Customer content by providing the Infrastructure.

9.2 Suitability. Customer is responsible for determining whether the Infrastructure is appropriate for Customer's security, compliance, export-control, and data-classification requirements.

9.3 Restricted Data. Unless otherwise agreed in writing, Customer shall not upload protected health information, payment card data, or other highly regulated data to the Infrastructure.

9.4 Data Return and Destruction. At the end of the Service Term or earlier termination of an Order, Customer is responsible for exporting Customer's data, models, software, configurations, and workload outputs from the Infrastructure before access ends. Upon Customer's written request made before the end of the Service Term or within five (5) days after termination, and to the extent technically practicable, Provider shall provide commercially reasonable assistance to return or enable export of such Customer data then remaining on the Infrastructure. After access ends and any requested return or export is completed, Provider shall delete or destroy Customer data remaining on the Infrastructure within thirty (30) days, except for archival copies retained in ordinary-course backups, security logs, or records required by law, which will remain protected under this Agreement until deleted in accordance with Provider's ordinary retention practices.

10. Confidentiality

10.1 Protection. Each Party may receive non-public business, technical, pricing, security, or operational information of the other Party. The receiving Party shall protect such information using reasonable care and shall use it only to perform or receive Services under this Agreement.

10.2 Exclusions. Confidentiality obligations do not apply to information that is independently developed, publicly available without breach, lawfully received from a third party, or already known without restriction.

11. Liability and Indemnification

11.1 Exclusion of Indirect Damages. In no event shall either Party be liable for indirect, incidental, special, consequential, exemplary, or punitive damages, or for lost profits, lost revenue, loss of data, or business interruption, arising out of or related to this Agreement.

11.2 Liability Cap. Each Party's aggregate liability arising out of or related to this Agreement and all Orders hereunder, regardless of the form or theory of action, shall not exceed two hundred fifty thousand dollars (\$250,000). The foregoing limitation shall not apply to: (a) Customer's obligation to pay all fees due or to become due under each Order in accordance with this Agreement, including fees for committed reserved capacity under Section 6.2; (b) a Party's fraud or willful misconduct; or (c) liability that cannot

be limited under applicable law. Service Credits under Schedule 1, including cash payments of unused Service Credits under Section 8.3 of Schedule 1, are not subject to, and shall neither count toward nor reduce, the foregoing aggregate liability cap. Such Service Credits remain subject to the applicable monthly cap in Section 8.2 of Schedule 1. Provider's obligation to refund prepaid fees under Section 7.5 is not subject to, and shall neither count toward nor reduce, the foregoing aggregate liability cap.

11.3 Provider IP Indemnity. Provider shall defend, indemnify, and hold harmless Customer and its officers, directors, employees, and affiliates from and against any third-party claim alleging that Customer's authorized use of the Infrastructure as provided by Provider under this Agreement infringes or misappropriates that third party's intellectual property rights. This obligation does not apply to claims arising from Customer's workloads, data, models, software, or instructions; combinations with items not provided by Provider; modifications not made by Provider; use outside the scope of this Agreement; or Customer's breach of this Agreement. If such a claim is made or, in Provider's reasonable opinion, is likely to be made, Provider may procure the right for Customer to continue using the affected Infrastructure, replace or modify the affected component so it is non-infringing and materially equivalent, or terminate the affected Services and refund prepaid fees for unused reserved capacity under Section 7.5.

11.4 Mutual Third-Party Claim Indemnity. Each Party shall defend, indemnify, and hold harmless the other Party and its officers, directors, employees, and affiliates from and against any third-party claim, damages, losses, liabilities, costs, and expenses, including reasonable attorneys' fees, arising out of or relating to the indemnifying Party's breach of this Agreement, gross negligence, or willful misconduct.

11.5 Indemnification Process. The indemnified Party shall promptly notify the indemnifying Party of any claim for which indemnification is sought, provide reasonable cooperation, and allow the indemnifying Party to control the defense and settlement of the claim. The indemnifying Party may not settle any claim in a manner that admits fault by the indemnified Party, imposes non-monetary obligations on the indemnified Party, or requires payment by the indemnified Party without the indemnified Party's prior written consent, not to be unreasonably withheld, conditioned, or delayed.

12. Sublease on the Ornn Exchange.

Customers may sublease all or any part of the Infrastructure to third parties for any period within the Service Term by listing it on the Ornn exchange, subject to the Ornn terms applicable to such listings. Customers shall not sublease or otherwise make the Infrastructure available to third parties by any other means without Provider's prior written consent.

13. Miscellaneous

13.1 Force Majeure. Neither Party shall be liable for delay or failure to perform due to events beyond its reasonable control, including acts of God, natural disasters, war, terrorism, civil unrest, labor disruptions, governmental actions, power or telecommunications failures, supply-chain failures, or data center, network, or cloud-provider outages not caused by the affected Party. The affected Party shall give prompt written notice describing the event, its expected duration, and the obligations affected, and shall use commercially reasonable efforts to mitigate the effects of the event and resume performance. If a force majeure event continues for more than thirty (30) consecutive days and materially prevents performance, either Party may terminate the affected Order upon written notice, with any refund of unused prepaid reserved capacity handled under Section 7.5.

13.2 Amendment. This MSA may be amended only by a written instrument signed by both Parties. An Order may be amended only by a written instrument signed by both Parties referencing that Order.

13.3 Governing Law. This Agreement shall be governed by the laws of the State of Delaware, without regard to conflict-of-law rules.

13.4 Entire Agreement. This MSA, including Schedule 1 (SLA) and each executed Order (including its exhibits), constitutes the entire agreement between the Parties with respect to its subject matter and supersedes all prior or contemporaneous agreements and understandings on that subject matter.

13.5 Assignment. Neither Party may assign or transfer this Agreement, in whole or in part, without the other Party's prior written consent, not to be unreasonably withheld, conditioned, or delayed. Either Party may assign this Agreement without consent to an affiliate or in connection with a merger, reorganization, change of control, or sale of all or substantially all of its assets or business to which this Agreement relates, provided that the assignee assumes the assigning Party's obligations and the assigning Party gives the other Party written notice. Any assignment in violation of this Section is void. No permitted assignment relieves either Party of obligations accrued before the assignment.

13.6 Notices. All formal notices under this Agreement must be in writing and delivered by personal delivery, nationally recognized overnight courier, registered or certified mail with return receipt requested, or email with confirmation of transmission; provided that notices of breach, termination, indemnity claims, or legal proceedings sent by email must also be sent by courier or registered or certified mail. Notices are effective upon personal delivery, upon confirmed delivery by courier or mail, or, for permitted email notices, on the next business day after transmission if no bounceback or delivery failure is received. Notices to Customer must be sent to Ornn Clearing LLC, 447 Broadway FL4, New York City, NY 10013, Attn: Legal Notices, with an email copy to legal@ornn.com. Notices to Provider must be sent to the notice address and email set forth on the signature page or the applicable Order. Either Party may update its notice address by giving notice under this Section.

13.7 Publicity. Neither Party may use the other Party's name, trademarks, service marks, logos, or other brand identifiers in any press release, customer list, case study, marketing material, public announcement, or other public communication without the other Party's prior written consent, except to the extent disclosure is required by applicable law or stock exchange rule.

13.8 Survival. Sections 6 (Financial Terms) with respect to accrued payment obligations, unpaid amounts, and refund obligations; 9.1 and 9.4 (Customer ownership and data return/destruction); 10 (Confidentiality); 11 (Liability and Indemnification); 13.3 (Governing Law); 13.6 (Notices); this Section 13.8; and any other provision that by its nature should survive, survive expiration or termination of this Agreement.

14. Signatures

The Parties have executed this Master Services Agreement as of the Effective Date.

CUSTOMER: ORNN CLEARING LLC, a Delaware limited liability company

By: _____

Name: Jack Minor

Title: COO

Notice Address: 447 Broadway FL4, New York City, NY 10013

Notice Email: jack@ornn.com

Date: _____

PROVIDER: _____ ,

By: _____

Name: _____

Title: _____

Date: _____

Notice Address: _____

Notice Email: _____

SCHEDULE 1

Service Level Agreement and Working Node Standard

1. Incorporation and Scope

This Service Level Agreement and Working Node Standard (this “SLA”) is incorporated into and forms part of the Master Services Agreement between Ornn Clearing LLC (“Ornn” or “Customer”) and the Provider identified in that agreement, and applies to each Order. Capitalized terms used but not defined in this SLA have the meanings given to them in the MSA or the applicable Order.

This SLA applies to all GPU compute capacity, nodes, clusters, networking, storage, management access, support, monitoring, and related services that Provider makes available to Customer under an Order (the “Services”).

If this SLA conflicts with the MSA or an Order, the order of precedence in Section 1.4 of the MSA controls. This SLA controls with respect to service levels, availability measurement, working-node requirements, downtime remedies, monitoring, and support obligations, unless an Order expressly states that it modifies this SLA for that Order.

2. Availability Commitment

Provider shall use commercially reasonable efforts to provide the Services with at least 99.9% Availability during each calendar month of the Service Term of each Order, as measured through monitoring, Customer reports, Provider notices, logs, and technical validation. The Availability Commitment applies to the ordered Services as a whole unless the applicable Order states a more specific node-level or cluster-level commitment.

A node is available while it materially satisfies the Working Node Standard. A cluster, pool, or reservation is available to the extent the relevant ordered capacity is materially usable for the purpose, SKU, topology, and performance level represented in the Order. Minor, temporary, or non-customer-impacting deviations do not by themselves constitute Downtime.

Provider shall deliver each node in a condition that materially satisfies the Working Node Standard at the start of the applicable Order and shall use commercially reasonable efforts to maintain that condition throughout the Service Term.

3. Working Node Requirement

A “Working Node” means a node that materially satisfies the baseline functionality, performance, interconnect, health, management, and validation requirements in Appendix A. A node that materially fails Appendix A will be treated as not satisfying the Working Node Standard only for the period of material failure, unless continued use would create customer-impacting risk.

The Working Node Standard is self-contained and does not require reference to any external third-party document. The Parties may update Appendix A from time to time by written agreement to reflect new SKUs, vendor-recommended diagnostics, security requirements, or operational lessons learned; provided that any material update will apply to an existing Order only after reasonable notice and a commercially reasonable implementation period, unless the Parties agree otherwise.

4. Downtime

“Downtime” means a material disruption, impairment, failure, or unavailability lasting five (5) or more consecutive minutes and involving any of the following:

- all or a material part of the ordered Services failing to materially satisfy the Working Node Standard;
- a node, cluster, interconnect, storage subsystem, management interface, or related service not materially functioning or not materially performing as represented in the MSA or Order;
- a node failing to respond to Customer’s monitoring, provisioning, health-check, SSH, API, or other operational tooling in a manner that materially impairs use of the Services;
- material loss, degradation, throttling, misconfiguration, or driver/runtime failure affecting GPUs, CPUs, RAM, local storage, NVLink, InfiniBand, Ethernet, power, cooling, kernel, firmware, CUDA, NCCL, container runtime, identity, or access control; or
- any other condition that materially prevents or materially degrades use of the Services by Customer.

A “Downtime Period” means any period comprising five (5) or more consecutive minutes of Downtime. Unless a more precise period is determined from reliable telemetry, a Downtime Period begins at the earliest validated point of material customer impact and ends when the affected Services are materially restored. If a condition constituting Unexcused Downtime renders a cluster, pool, or other ordered capacity materially unusable for its ordered purpose, every node rendered materially unusable by that condition shall count toward Node Downtime Minutes for the period of that impact, even if the underlying fault is confined to one node or a shared component.

5. Unexcused Downtime

“Unexcused Downtime” means Downtime that is not Excused Downtime. Provider is responsible for Unexcused Downtime to the extent caused by Provider-controlled infrastructure, operations, subcontractors, or Provider-managed dependencies.

For each calendar month in which Availability falls below the applicable commitment due to Unexcused Downtime, Customer shall receive the applicable service credit calculated under Section 8. Provider shall also perform the support, remediation, and reporting obligations in this SLA.

6. Excused Downtime and Maintenance

“Excused Downtime” means Downtime caused by Maintenance that satisfies this Section 6 or by a cause expressly excluded under Section 7. Provider shall provide commercially reasonable documentation supporting any claimed Excused Downtime.

6.1. Scheduled Maintenance

Scheduled maintenance is excused if Provider gives Customer commercially reasonable prior written notice, with a target of at least twenty-four (24) hours where practicable, describes the expected scope and duration, uses commercially reasonable efforts to limit customer impact, and restores all affected capacity promptly after completion.

6.2. Emergency Maintenance

Emergency maintenance is excused to the extent it is reasonable and necessary to prevent material harm to infrastructure, data, security, safety, or customers. Provider shall give Customer prompt notice and commercially reasonable follow-up detail regarding cause, impact, mitigation, and recurrence risk.

6.3. Effect of Excused Downtime

During Excused Downtime, affected Services may be suspended for the period reasonably necessary to perform the maintenance. For reserved or committed capacity, the Parties will discuss in good faith an extension or other operational accommodation if Excused Downtime materially exceeds the scheduled or expected window.

7. Exclusions

Downtime will not count as Unexcused Downtime to the extent caused by: (a) Customer's unauthorized modification of the node outside the agreed operating model; (b) Customer workloads that exceed documented limits or violate the Agreement after Provider has followed the required escalation process; (c) Customer systems that are outside Provider's control and not necessary for Provider to deliver the ordered capacity; (d) scheduled or emergency maintenance; (e) public internet, carrier, power, facility, or cloud events outside Provider's reasonable control; or (f) force majeure events, in each case only for the period during which Provider could not reasonably prevent or mitigate the impact.

The exclusions do not apply to the extent the claimed excluded cause resulted primarily from Provider's failure to maintain commercially reasonable power, cooling, configuration, firmware, drivers, monitoring, documented operational procedures, security controls, or recovery processes.

8. Service Credits

8.1. Downtime Percentage

Availability and Service Credits shall be measured separately for each Order, or for a separately priced cluster or other capacity only if the Order expressly identifies it as a separate measurement unit and allocates its fees. For each calendar month, the "Downtime Percentage" equals aggregate Node Downtime Minutes for that same capacity divided by aggregate ordered node-minutes for that capacity during the Service Term in that month, expressed as a percentage. Ordered node-minutes are the sum of the number of ordered nodes for each minute of that period, including minutes of Downtime and excluding the Testing Period. This calculation reflects partial months and agreed changes in node count. Monthly "Availability" equals one hundred percent (100%) minus the Downtime Percentage. A "Node Downtime Minute" is one of the ordered node-minutes described above during which the relevant node experiences Unexcused Downtime, including under Section 4. Each node-minute is counted only once, regardless of concurrent failures.

8.2. Monthly Service Credit

If monthly Availability falls below 99.9% due to Unexcused Downtime, the "Monthly Service Credit" equals the fees paid or payable for the same capacity and monthly period measured under Section 8.1, before deduction of Service Credits or disputed amounts and including prepaid fees allocated to that period, multiplied by the applicable credit percentage: (a) ten percent (10%) where Availability is at least 99.0% but below 99.9%; (b) thirty percent (30%) where Availability is at least 95.0% but below 99.0%; or (c) one hundred percent (100%) where Availability is below 95.0%. No Service Credit is due where

Availability is at least 99.9%. Use unrounded Availability to select the single applicable tier; tiers are not cumulative. No further multiplication by the Downtime Percentage applies. Aggregate Service Credits for a measurement unit and calendar month shall not exceed one hundred percent (100%) of those monthly fees, excluding taxes.

8.3. Application of Service Credits

Service Credits apply against the next invoice or payment otherwise due for the affected Services. Any unused balance carries forward without expiry. Upon expiration or termination of an Order, Provider shall pay Customer any unused, undisputed Service Credits attributable to that Order within thirty (30) days. For a timely claim or calculation dispute resolved after expiration or termination, Provider shall pay any additional unused Service Credits within thirty (30) days after determination. The rights and obligations under Sections 8.3, 8.4, and 15 survive expiration or termination, including the right to submit a timely claim after the Order ends. A dispute shall not delay application or payment of any undisputed credit.

8.4. Credit Process

Provider shall calculate and issue Service Credits for Unexcused Downtime established by its records or other reliable evidence without requiring a separate Customer request. Customer may request additional or corrected Service Credits within sixty (60) days after the end of the month in which the Unexcused Downtime occurred and shall provide reasonable supporting information, which may include Customer monitoring and logs. Provider shall review the request promptly and in good faith and apply or pay any undisputed credit in accordance with Section 8.3.

8.5. Limited SLA Remedy

Service credits are Customer's sole and exclusive remedy for ordinary Availability failures under this SLA. This limitation does not apply to fraud, willful misconduct, gross negligence, confidentiality breaches, security incidents, data loss, indemnity claims, intentional misuse of capacity, or any claim for which exclusive-remedy treatment would be prohibited by the Agreement or applicable law. This limitation does not restrict Provider's repair or replacement obligations under this SLA or Customer's termination or refund rights under the MSA.

9. No Use of Non-Working Nodes

Provider shall not knowingly route Customer production workloads to capacity experiencing material Downtime during the Service Term of the applicable Order unless Customer has approved continued use, accepted partial capacity, released the capacity, or the Parties agree in writing. Provider may use affected capacity for repair, testing, diagnostics, and restoration. If a replacement or release mechanism is available, Provider shall coordinate with Customer in good faith and shall not represent materially impaired capacity as production-ready until Provider has reasonably validated restoration.

10. Monitoring, Access, and Evidence

Provider shall permit Customer to use monitoring agents, scripts, APIs, telemetry collectors, logs, access, out-of-band checks, and related tools reasonably necessary to measure Availability, detect Downtime, and validate Working Node status, subject to reasonable security, access, and change-management controls.

Provider shall maintain commercially reasonable operational records for the Services, including maintenance windows, incident timelines, facility events, network events, hardware replacements, firmware and driver changes, diagnostic results, and customer-impact summaries. Provider shall provide relevant records to Customer upon reasonable request and shall preserve records relating to any material incident for at least one (1) year.

11. Support Services

11.1. Customer First-Tier Responsibilities

Customer will ordinarily provide the first level of support to its own users and workloads for triage, account issues, and Customer-managed software, and will determine whether an issue requires Provider support based on monitoring, reports, logs, or reasonable operational judgment.

11.2. Provider Support

Except for Customer-managed software and workloads, Provider shall provide all technical support reasonably necessary for the Services to perform in accordance with the MSA, this SLA, and each Order.

11.3. Severity Levels and Response Times

Provider shall use commercially reasonable efforts to meet the support response targets in Appendix C. Provider shall maintain escalation contacts capable of acting on critical incidents for production Services unless an Order expressly states a different support model.

12. Incident Response and Root Cause Analysis

For any Severity 1 or repeated Severity 2 incident, Provider shall commercially reasonably investigate, mitigate, restore service, preserve relevant evidence, and keep Customer informed at a reasonable frequency given the incident severity.

13. Replacement, Cure, and Chronic Failure

“Chronic Failure” means a node experiencing three (3) or more separate periods of Unexcused Downtime within any rolling thirty (30)-day period caused by faults in that node’s hardware, firmware, or Provider-managed software, with restoration to the Working Node Standard between incidents. Customer may reasonably request that Provider replace, remove, quarantine, or revalidate any node that materially fails the Working Node Standard, experiences repeated Downtime, exhibits intermittent errors, materially underperforms the accepted hardware profile, or presents a material security, reliability, facility, or customer-impact risk. Provider may propose a commercially reasonable cure plan before replacement unless immediate replacement is necessary to prevent material customer impact. Notwithstanding the foregoing, upon Customer’s written request following Chronic Failure of a node, Provider shall promptly replace that node with equivalent or better capacity that satisfies the applicable Order and the Working Node Standard, at no additional charge to Customer. The Chronic Failure threshold does not limit Customer’s other rights under this Section.

14. Security and Operational Controls

Provider shall maintain commercially reasonable administrative, technical, physical, and operational safeguards for the Services, including secure access control, least-privilege administration, patch and vulnerability management, logging, physical security, power and cooling safeguards, network segmentation, incident response procedures, and controls reasonably designed to prevent unauthorized customer access or cross-tenant exposure.

Provider shall notify Customer promptly of any actual or suspected security incident, unauthorized access, data exposure, compromise of credentials, supply-chain compromise, firmware compromise, or facility event that may affect the Services or Customer.

15. Monthly Reporting and Disputes

Either Party shall raise any good-faith dispute regarding a service credit calculation within thirty (30) days after receiving the calculation. The Parties shall cooperate to resolve the dispute promptly, and any

undisputed credit will be applied or paid in accordance with Section 8.3. This Section 15 shall not shorten Customer's claim period under Section 8.4.

16. Contact

Questions and operational escalations to Provider regarding this SLA should be sent to the Provider contacts identified in the applicable Order or, if none are identified, to Provider's notice email on the signature page. Customer's operational contact is legal@ornn.com, unless the Order identifies another Customer contact. Formal notices under this SLA shall be given in accordance with Section 13.6 of the MSA.

SLA Defined Terms

Defined Term	Meaning
Order	An order form executed under the MSA under which Provider supplies Services to Customer.
Availability	The monthly percentage calculated under Section 8.1, measured against the commitment in Section 2 or the applicable Order.
Customer First-Tier Responsibilities	The first level of support performed by Customer for its own users and workloads, and the determination of whether Provider support is required.
Downtime	The meaning given in Section 4.
Downtime Period	The meaning given in Section 4.
Excused Downtime	Downtime excluded under Section 6 or Section 7.
Maintenance	Scheduled or emergency maintenance that substantially satisfies Section 6.
Node Downtime Minute	The meaning given in Section 8.1, including capacity rendered unusable under Section 4.
Service Credit	The service credit calculated under Section 8 for Unexcused Downtime during the applicable calendar month.
Services	The GPU compute capacity, nodes, clusters, networking, storage, management access, support, monitoring, and

	related services supplied under an Order.
Unexcused Downtime	Downtime that is not Excused Downtime.
Working Node	A node materially satisfying Appendix A and applicable requirements in the MSA and Order.

APPENDIX A

Working Node Standard

A node satisfies the Working Node Standard only if:

- a. All GPUs, CPUs, system memory, storage devices, NICs, HCAs, interfaces, and other material components specified in the applicable Order are present and visible.
- b. GPU quantity, model, serial numbers, and UUIDs materially match the applicable Order.
- c. HCA quantity, model, port state, link width, and link rate materially match the applicable Order.
- d. PCIe, NVLink, NVSwitch, InfiniBand, Ethernet, and other material interconnects required by the Order are present and operating without material errors.
- e. The operating system, NVIDIA driver, CUDA, NCCL, PyTorch, DCGM, RDMA utilities, and other tooling required for validation and Customer use are installed and operational.
- f. The node is reachable through the agreed production and management access paths.
- g. No unauthorized workload, unresolved material hardware fault, material thermal issue, or material network fault is present.
- h. The node passes all applicable validation requirements in this Appendix.

A passing node result does not establish that a cluster has passed. Where the Order includes a multi-node cluster, the complete selected cluster must also pass the multi-node validation in Section A.2.4.

A.2. Pre-Turnover Validation

Provider shall validate every node before turnover to Customer.

Provider shall use one of the following validation profiles:

Standard. Standard is the routine pre-turnover profile. It includes hardware inventory and initial health checks, one complete NVIDIA DCGM Level 3 diagnostic with basic GPU memory enabled, limited host-memory pattern testing, local GPU communication testing, forced InfiniBand or RDMA-path communication testing, final health checks, and retention of test evidence.

Deep. Deep includes every Standard requirement and additionally performs all eleven extended DCGM GPU memory patterns and a full host-memory pattern pass. Deep is selected expressly for extended

qualification or investigation. Deep does not begin automatically because of enrollment, a warning, or an incomplete Standard result.

Standard and Deep validation do not include a separate duration-based GPU burn-in, GEMM benchmark, or training workload. Any separately requested burn-in does not replace the validation required by this Appendix.

For multi-node turnover, the validation scope shall include every node and GPU in the ordered cluster and the combined test in Section A.2.4. Individual node passes do not establish a cluster pass.

Before starting validation, Provider shall select the applicable nodes, profile, cluster scope, coordinator information, and approved performance reference. Provider shall confirm that the selected nodes are available for testing.

Tests proceed sequentially on each node. A failed or incomplete required check stops subsequent stress stages, although final health and evidence collection should still be attempted.

A required test that is missing, skipped, unavailable, incomplete, cancelled, or unsupported does not pass.

A.2.1. GPU Health and Compute Diagnostics

For both Standard and Deep validation, Provider shall run one complete NVIDIA DCGM Level 3 diagnostic on every node with the basic GPU memory test expressly enabled:

`dcgmi diag -r 3 -j --parameters memory.is_allowed=true`

The validation runner shall verify per-GPU results for:

- a. Software and driver health.
- b. Basic GPU memory allocation, access, and integrity.
- c. PCIe communication.
- d. Diagnostic computation.
- e. GPU memory bandwidth.
- f. Targeted GPU stress.
- g. Targeted GPU power.
- h. All additional subtests reported by DCGM.

No required DCGM subtest may be disabled to shorten the validation.

Provider may check the installed DCGM CLI and host-engine versions against NVIDIA's applicable stable package channel. A newer available version or an unavailable version lookup may be recorded as a warning, but does not replace or override the diagnostic result. Validation shall use the installed DCGM version and shall not automatically upgrade DCGM or restart its services.

The diagnostic passes only if:

- a. The diagnostic completes successfully.

- b. Every required result is present for every GPU.
- c. Every required DCGM subtest passes.
- d. No GPU disappears during validation.
- e. No material XID error occurs.
- f. No uncorrectable ECC error occurs.
- g. No pending or failed row remapping is detected.
- h. No hardware thermal slowdown occurs.
- i. No external power-brake condition occurs.

Normal software power-cap management is not, by itself, a hardware failure.

Provider shall retain the DCGM output, before-and-after health evidence, and GPU health samples collected during validation.

A.2.2. GPU and Host Memory Diagnostics

Standard Profile

Under Standard validation, GPU memory coverage is provided by the basic DCGM memory test in Section A.2.1 and is not duplicated.

Provider shall also run a limited host-memory pattern test covering:

- a. Stuck Address.
- b. Random Value.
- c. Compare XOR.

Deep Profile

Under Deep validation, Provider shall additionally run DCGM GPU memtest patterns zero through ten as eleven separate invocations.

Each invocation shall:

- a. Select one specific memory pattern.
- b. Request sixty seconds of testing for that pattern.
- c. Include every GPU.
- d. Produce a passing result for every GPU.

A combined or unverified memtest command does not establish that all eleven patterns were completed.

Deep validation shall also run one full default host-memory pattern pass.

Host-Memory Coverage

For both profiles, concurrent memory-test workers shall collectively request at least ninety percent of MemAvailable, measured immediately before the host-memory test.

Every worker must:

- a. Lock its requested allocation.

- b. Complete every required pattern.
- c. Report the actual tested allocation.
- d. Exit successfully.

The memory validation passes only if:

- a. All selected GPU and host-memory tests complete.
- b. No comparison error occurs.
- c. No new uncorrectable ECC error occurs.
- d. No pending or failed row remapping is detected.
- e. No material host hardware error is detected.

An allocation shortfall, unfinished worker, missing pattern result, skipped GPU, or missing evidence produces an Incomplete result rather than a Pass.

A.2.3. Single-Node GPU Communication

For both Standard and Deep validation, Provider shall run the versioned Ornn NCCL v1 all-reduce harness through PyTorch with one process per GPU.

This harness is not the NVIDIA nccl-tests executable. Performance references from another harness are not interchangeable unless the Parties approve the substitution in writing.

Each communication test shall use:

- a. Payloads of 8 bytes, 1 MiB, 1 GiB, and 16 GiB.
- b. Two warmup repetitions.
- c. Five measured repetitions.
- d. One process for each participating GPU.
- e. A correctness check of every output element after every reduction.
- f. The median timing of the slowest rank for each payload.

The harness shall report algorithm bandwidth and all-reduce bus bandwidth. Bus bandwidth shall be calculated as algorithm bandwidth multiplied by:

$$2 \times (\text{number of ranks} - 1) \div \text{number of ranks}.$$

Local GPU Path

All local GPUs shall participate. NCCL logs must show use of the expected local P2P or NVLS path. Provider shall retain the GPU topology used to interpret the result.

NIC Path

Provider shall run a separate NIC-path test with local peer-to-peer, shared-memory, and NVLS communication disabled and NCCL's InfiniBand transport forced:

NCCL_P2P_DISABLE=1

NCCL_SHM_DISABLE=1

NCCL_NVLS_ENABLE=0

NCCL_NET=IB

NCCL transport logs and before-and-after HCA port counters must demonstrate RDMA traffic on the participating active ports without an increase in material link-error counters.

Configuration flags alone do not prove that the required transport was used. A local NIC-path result does not independently certify leaf-switch routing or whole-cluster performance.

The single-node communication validation passes only if:

- a. Every GPU rank completes every payload.
- b. Every reduction produces the correct result.
- c. The required local and NIC transports are demonstrated.
- d. Required HCA traffic is observed.
- e. Material HCA error counters do not increase.
- f. The 16 GiB bus bandwidth for each required path meets the approved reference in Section A.3.

Missing tools, missing ranks, timeouts, incomplete measurements, unverified transport, or a missing performance reference produce an Incomplete result. Incorrect reductions, increased material port-error counters, or bandwidth below the approved reference produce a Fail result.

A.2.4. Multi-Node GPU Communication

For a multi-node cluster, Provider shall perform a combined Ornn NCCL v1 all-reduce test across every selected node and GPU after the required node-level checks have completed successfully.

Before the combined test, Provider shall freeze:

- a. Cluster membership.
- b. Node and rank ordering.
- c. Coordinator address and port.
- d. Validation profile.
- e. Applicable performance reference.

Participating nodes must use materially matching:

- a. GPU models and GPU counts.
- b. NVIDIA drivers.
- c. PyTorch versions.
- d. CUDA versions.
- e. NCCL versions.
- f. Ornn NCCL harness versions.

The coordinator address must belong to the selected coordinator node and must be reachable through existing peer connectivity.

Provider shall not distribute SSH keys or change firewall rules to conduct the test. Each enrolled node shall launch only its local GPU processes, and existing node-to-node connectivity shall carry the collective.

The combined test shall use the same payloads, warmups, measured repetitions, correctness checks, timing method, and bandwidth calculation described in Section A.2.3.

The cluster validation passes only if:

- a. Every selected node and GPU participates.
- b. Every member completes the combined test.
- c. Every reduction produces the correct result.
- d. The required InfiniBand or RDMA transport is demonstrated.
- e. No node or GPU drops from the collective.
- f. No NCCL, CUDA, RDMA, transport, timeout, or fabric error prevents completion.
- g. The cluster's 16 GiB bus bandwidth meets the approved cluster reference in Section A.3.

A missing, failed, timed-out, cancelled, or unfinished member prevents a cluster Pass.

Where Provider performs a measurement-only run to establish an initial reference, the measurements may be retained, but the cluster result remains Incomplete until an approved reference is recorded and satisfied.

A.3. Configuration-Specific Performance References

Before a qualification result may pass, Provider shall record an approved, named performance reference containing separate minimums for:

- a. Local GPU-path bandwidth.
- b. NIC-path bandwidth.
- c. Whole-cluster bandwidth, where the Order includes a multi-node cluster.

The reference must materially match:

- a. GPU model and count.
- b. Node and GPU topology.
- c. HCA and link configuration.
- d. Driver, PyTorch, CUDA, and NCCL versions.
- e. Harness and payload.
- f. Cluster size, where applicable.

A reference may be based on reviewed measurements from known-healthy equipment with the same material configuration or on a requirement separately agreed by the Parties.

If the Parties adopt a percentage-of-baseline policy, the resulting absolute minimum in GB/s and its derivation shall be retained with the reference.

A median derived from unqualified machines does not establish a healthy baseline. A performance reference may not be selected or reduced after reviewing a failing test result.

Bandwidth under this Appendix is reported in decimal GB/s using the 16 GiB payload. The reported all-reduce bus-bandwidth metric is not a direct NIC line-rate measurement.

No fixed bandwidth threshold applies universally to every SKU and topology. If an approved reference has not been supplied, Fabric may retain the measurements, but the result is Incomplete rather than Pass.

A.4. Results, Evidence, and Retesting

Pass means every required check completed with passing evidence and every applicable performance reference was met.

Fail means a measured health, correctness, or performance requirement failed.

Incomplete means a required test, tool, reference, rank, sample, result, or evidence item is missing, skipped, unavailable, unsupported, cancelled, or unfinished.

A successful process exit alone is insufficient to establish a Pass. Cancellation cannot produce a Pass. Execution deadlines are maximum safety limits and are not minimum test durations.

Provider shall capture and retain commercially reasonable records for each validation, including:

- a. Validation profile and version.
- b. Node and cluster membership.
- c. Start and completion timestamps.
- d. DCGM JSON and per-GPU results.
- e. Selected GPU memory patterns and results.
- f. Host-memory worker allocations, logs, patterns, and outcomes.
- g. Collective timing, bandwidth, correctness, and transport results.
- h. Before-and-after GPU, kernel, and HCA error evidence.
- i. NVIDIA driver, CUDA, NCCL, PyTorch, and DCGM versions.
- j. GPU inventory and serial numbers.
- k. HCA and port identifiers.

Fabric shall retain validation configuration, state, and reports in its benchmark data store and shall save result and log artifacts in the applicable benchmark storage bucket.

Missing or undeliverable evidence prevents a Pass. Evidence may be compressed in transit, but must be restored for storage and review.

Provider shall provide a summary of the applicable validation results to Customer at turnover.

Customer may reasonably request that the Provider rerun a test using Customer tooling where a material issue is suspected.

If a node materially fails a required validation test, the affected capacity does not satisfy the Working Node Standard unless Customer accepts the deviation or partial capacity in writing.

Provider may return failed capacity to service only after correcting the failure and obtaining a passing result from a new validation run. Earlier evidence and failure records shall remain unchanged.

A.5. Amendments, Substitutions, and Exceptions

A SKU-specific substitute requires written approval and must measure the same health, correctness, performance, and usability condition without materially weakening this Appendix.

Provider may modify or waive a validation requirement only when:

- a. The Parties agree in writing to a different testing schedule, duration, tool, or procedure for the applicable node or cluster.
 - b. The identical SKU and material configuration successfully passed the applicable validation requirements within the preceding seventy-two hours.
 - c. Customer agrees in writing to waive pre-turnover validation for an emergency replacement or hot swap.
- Any approved exception must identify the affected equipment, omitted or modified requirement, operational impact, mitigation, and required remediation date.
- Customer acceptance during the Testing Period remains governed by Section 4.4 of the MSA.

APPENDIX B

Service Credit Examples

The following examples are included to make the service credit mechanics clear. They do not modify the applicable credit cap or any express terms in the MSA or an Order. The table assumes sixteen (16) nodes ordered for a full thirty-day month (43,200 minutes), with only the stated Unexcused Downtime. Displayed Availability is rounded for illustration; tier selection uses the unrounded value.

Scenario	Calculation	Result
One node down for 30 minutes in a 16-node order	30 Node Downtime Minutes / (16 nodes x 43,200 minutes)	99.9% Availability; no Service Credit.
Four nodes down for 15 minutes in a 16-node order	60 Node Downtime Minutes / (16 nodes x 43,200 minutes)	99.9% Availability; no Service Credit.
Eight nodes down for a full hour in a 16-node order	480 Node Downtime Minutes / (16 nodes x 43,200 minutes)	99.9% Availability; no Service Credit.
All nodes down for a full hour	960 Node Downtime Minutes / (16 nodes x 43,200 minutes)	99.8% Availability; 10% Service Credit.

Example credit: if monthly fees for the capacity measured under Section 8.1 are \$10,000.00 and the Downtime Percentage for the month is 2%, Availability is 98% and the Monthly Service Credit is \$3,000.00 (30% of \$10,000.00), subject to the 100% monthly cap and the credit process in Section 8. If Availability is below 95%, the Monthly Service Credit is \$10,000.00 (100% of those monthly fees).

APPENDIX C

Support Severity Levels

Severity	Description	Acknowledgement Target	Provider Obligation
Severity 1	Complete service outage; cluster or reserved pool unavailable; security incident; data exposure risk; widespread customer impact; failure that prevents material customer workloads.	1 hour, 24x7	Commercially reasonable continuous work until mitigated; updates every 2 hours during active impact; root cause summary within 10 business days upon request.
Severity 2	Material degradation; individual node failure affecting active or reserved capacity; repeated job failures; degraded network or storage; customer-impacting performance below accepted thresholds.	4 hours, 24x7 for production Services	Commercially reasonable active work until mitigated; daily updates during active impact; root cause summary if repeated or requested.
Severity 3	Non-critical issue, partial impairment, validation warning, support question, or issue with workaround that does not materially prevent customer workloads.	1 business day	Commercially reasonable work during business hours; periodic updates while open.
Severity 4	Documentation request, planned change coordination, non-urgent improvement, or routine operational request.	2 business days	Handle through normal operational queue and agreed change process.

Provider shall maintain current escalation contacts for Severity 1 and Severity 2 incidents, including names or roles, email addresses, phone numbers, and escalation paths. Provider shall give Customer reasonable notice when changing any critical escalation contact.