



# Hyve Orion™


## Featuring NVIDIA HGX™ Rubin NVL8 Direct Liquid Cooled (DLC)

Built for AI and high-performance computing for every data center

Built on NVIDIA HGX Rubin, integrating eight GPUs with sixth-generation NVLink and a flexible host CPU architecture for large-scale AI training and inference, with high-bandwidth GPU connectivity and support for scale-up systems and clustered data center deployments.

### Key Features

- NVIDIA HGX™ Rubin NVL8 with NVLink™ fabric and integrated ConnectX®-9 networking
- 2U MGX-based architecture with separate CPU and GPU sleds
- Direct liquid cooling optimized for >20kW system power density
- High-bandwidth east-west networking up to 800G per GPU
- Front I/O serviceability with modular, rack-scale deployment support



## System Specifications

<p><b>Form Factor</b></p> <p>19" 2U rackmount, MGX-compatible architecture (87 mm H × 448 mm W × ~800 mm D)</p>	<p><b>CPU</b></p> <p>Dual Intel® Xeon® 6700 series (Granite Rapids-SP class)</p> <ul style="list-style-type: none"> <li>• High core count, PCIe Gen5 host connectivity</li> </ul>	<p><b>Expansion / I/O</b></p> <ul style="list-style-type: none"> <li>• MGX Type B and Type C modular bays</li> <li>• PCIe Gen5/Gen6 capable infrastructure (forward compatible)</li> </ul>
<p><b>GPU Platform</b></p> <p>NVIDIA HGX™ Rubin NVL8 (Direct Liquid Cooled)</p> <ul style="list-style-type: none"> <li>• 8× SXM8 modules (16 Rubin GPU dies total)</li> <li>• NVLink™ fully connected fabric (NVLink™ 6.0)</li> <li>• Up to 2.3 kW per GPU (TGP)</li> </ul>	<p><b>Memory</b></p> <p>Up to 32× DDR5 DIMMs</p> <ul style="list-style-type: none"> <li>• Up to 6400 MT/s</li> <li>• Configurable up to multi-terabyte capacity</li> </ul>	<p><b>Cooling</b></p> <p>Direct Liquid Cooling (DLC) – fully integrated</p> <ul style="list-style-type: none"> <li>• Cold plates for GPUs, NVSwitch, and CX-9</li> <li>• Integrated manifold and OSFP cooling loop</li> <li>• Universal Quick Disconnects (UQDs)</li> </ul>
<p><b>System Power</b></p> <ul style="list-style-type: none"> <li>• 54V DC busbar input</li> <li>• ~24 kW typical system TDP (up to ~30 kW+ peak design envelope)</li> </ul>	<p><b>Networking (East-West)</b></p> <p>Integrated NVIDIA ConnectX®-9</p> <ul style="list-style-type: none"> <li>• 8× OSFP ports</li> <li>• Up to 800 Gb/s per GPU (InfiniBand XDR / Ethernet)</li> </ul>	<p><b>System Management</b></p> <ul style="list-style-type: none"> <li>• NVIDIA HGX Management Controller (HMC)</li> <li>• DC-SCM (AST2600 + Root of Trust)</li> <li>• Redfish-compatible management stack</li> </ul>
<p><b>Storage</b></p> <ul style="list-style-type: none"> <li>• OS: 2× M.2 NVMe</li> <li>• Data: Up to 8× E1.S NVMe (front serviceable)</li> </ul>	<p><b>Networking (North-South)</b></p> <ul style="list-style-type: none"> <li>• 1× NVIDIA BlueField®-4 DPU (PCIe Gen5 x16)</li> <li>• Secure multi-tenant and infrastructure offload capabilities</li> </ul>	<p><b>Serviceability</b></p> <ul style="list-style-type: none"> <li>• Front I/O serviceable design</li> <li>• Cold aisle service model</li> <li>• Modular MGX tray-based architecture</li> </ul>

sales@hyvesolutions.com