

UMIACS Data Center Node and Infrastructure Lifecycle Support Policy

v1.0

UMIACS will support hardware nodes in our data center for their usable lives, as determined by the UMIACS Technical Staff, with certain limitations outlined in this lifecycle support policy. Nodes are defined as compute or storage systems that are either standalone or integrated into a cluster. The age of a node or system, whether purchased or donated/gifted, will be determined by the start date of its initial shipped manufacturer warranty.

UMIACS has a limited amount of personnel time to support legacy equipment. The data center also has three other limiting factors of capacity; physical number of racks, power distribution, and temperature control. This policy helps provide some reasonable expectation for the timeline of support that can be provided for this infrastructure and the equipment that is supported within it.

The UMIACS Technical Staff defines a Principal Investigator (PI) as the faculty member that initially purchases the equipment or the named faculty (Custodian) member that is delegated or designated by the original PI(s) to have current authority over the equipment.

The UMIACS Technical Staff will provide timely information to PIs whose equipment will be transitioning under this support policy. PIs will have a minimum grace period of 60 days between any change to nodes unless a catastrophic failure event occurs with the hardware.

Support Categories

To ensure staff can best prioritize support, there are three different service support categories:

Active -> Extended -> Legacy

Active service means that the equipment has an active support contract associated with it. The UMIACS Technical Staff will work at our highest priority to try to resolve issues and return equipment to operational status. The exact time to resolution is highly dependent on the specific warranty service(s) purchased with the equipment. The PI(s) of this equipment will be notified if repairs that would require additional funds due to something happening that is outside the scope of the warranty. Use and allocation of equipment in active service is governed by the PI(s).

Extended service means that the equipment no longer has an active support contract associated with it, but it has been less than 7 years since the start of its initial shipped manufacturer warranty. The UMIACS Technical Staff will work at our normal priority to return

equipment to service. For failures that require replacement hardware, the UMIACS Technical Staff will notify and work with the PI(s) to identify the cost and then let them decide whether or not to continue with repair from funds they can allocate. The PI(s) will assume the risk for all repairs out of a warranty period and the UMIACS Technical Staff will perform a best effort repair. If hardware can not be repaired or will not be repaired based on the decision of the PI(s), then staff will make a best effort to allow the user(s) to extract data from the nodes and then deprecate and surplus the equipment. Use and allocation of equipment in extended service level is governed by the PI(s).

Legacy service means that the equipment no longer has an active support contract associated with it, and it has been more than 7 years since the start of the initial shipped manufacturer warranty. When a node or system in this service level fails, the UMIACS Technical Staff will attempt the lowest priority for repairs. If the equipment is deemed to require significant repair or additional parts, or the failure of the equipment has happened more than once, the UMIACS Technical Staff is delegated the authority to surplus equipment that is no longer capable of performing the desired function. Use of equipment in the legacy service level is governed by the UMIACS Technical Staff and staff is delegated the authority to make it available in a shared, non-preemptible pool of resources for members of the entire UMIACS community.

Physical Data Center Capacity

Equipment located in our data centers is required to be able to be placed into racks. Racks have two different dimensions that staff will support:

- Standard Rack: 600mm width, 1070mm depth, 1991mm height
- Deep Rack: 600mm width, 1200mm depth, 1991mm height

Racks will often have a support lifetime well over 10 years. However, the UMIACS Technical Staff reserves the right to inform PI(s) that replacement of a rack is required at the time that it is no longer serviceable due to a physical failure. Staff considers all racks to be in active support service due to the limited failure rates.

All equipment, whether purchased or donated/gifted, must include the hardware required to mount it into racks, be of reasonable quality (as deemed by the Director of Computing Facilities), and allow reasonable maintenance on both the outside and inside of the equipment. Generally UMIACS does not purchase or support refurbished/used equipment however this can be requested to be reviewed by the Director of Computing Facilities.

Alternative rack dimensions may be considered based on the project and funding, and will require explicit approval of the Director of Computing Facilities. Any rack and equipment outside these default specifications will need to be reviewed and approved at 2 year intervals after the shortest warranty period of the systems contained within.

Rack units (typically referred to as a U) are the measurement of how much space a node or system will take when it is installed inside a rack. For our dimensions, staff supports 42U racks

(which is the maximum amount of space available). However, due to power and cooling requirements staff may not be able to fully populate all units. Most nodes and systems are 1-5U in size.

UMIACS provides a limited set of rack units (6) that are a base support for an individual standalone faculty without further investment. This provides for a basic set of equipment to be stood up without incurring the cost of rack, power distribution and networking costs. For node or systems investments that go beyond this base support, the UMIACS Technical Staff will provide the cost to purchase rack units, and power distribution and networking equipment to the PI(s).

Power Distribution

The UMIACS Technical Staff maintains a limited set of power distribution units (PDUs) that manage the step down in voltage to a useful level of between 200-240V from our uninterruptible power systems (UPSes). These PDUs provide a number of breaker positions that subsequently provide an interface to pull whip(s) to a rack. These whips provide receptacles and are run to racks so that the cabinet distribution units (CDUs) can provide power to nodes and systems. These CDUs have both temperature and humidity sensors that measure the performance of the rack.

UMIACS will allocate redundant power in all racks by default. This ensures that the rack has power from two discrete PDU/UPS combinations. This allows the UMIACS Technical Staff to service one of the PDU units without impacting all equipment in the rack. Staff expects all nodes and systems to have redundant power supplies by default and will only accept non-redundant systems by approval of the Director of Computing Facilities.

PI(s) are responsible for both the purchase and maintenance costs associated with CDUs in racks they have purchased. These typically have an operational lifetime of 7-10 years, however the CDUs may only provide a warranty through 2-3 years. In these cases, re-investment may be required before their normal operational lifetime to ensure that power redundancy is maintained.

Temperature Control

The data center uses a set of computer room air conditioning (CRAC) units to maintain temperature control for equipment. Our current data center facilities utilizes forced air from under a raised pressurized floor to provide cooling and humidity controls.

The primary UMIACS data center facility is now more than 35 years old and the floor is difficult to fully seal and pressurize. Thus locality in distance to these units can impact performance. Equipment will be placed in an best-effort optimal configuration, however there is a limit to the total cooling capacity of the room. This may restrict the UMIACS Technical Staff from adding further equipment in the future.

All nodes and systems placed in the data center must be able to conform to a specific range of temperatures that will allow for optimal operation of the equipment.

Networking

Networking in our data centers are provided by three different UMIACS network infrastructures; core routing, production networks, and research networks.

UMIACS core routing mediates access to our external routing and edge networks in our supported buildings, production, and research networks. This facility is currently a 100Gb per second network and it is maintained and supported financially by UMIACS. If research instrumentation requires additional capacity or capabilities, staff will work with the PI(s) who may need to provide additional funding to expand this infrastructure.

UMIACS production networks are our base network for core services in our data centers. These are commodity networks with lower performance ceilings than our research data networks. This network infrastructure is currently a 1/10Gb per second network and it is maintained and supported financially by UMIACS. Staff offers each UMIACS faculty up to 6 network drops that are serviced by this network infrastructure; however there is no guarantee on exact speed. PI(s) may make investments in top of rack (TOR) switches that are serviced by this network that expand the number of network drops on a per rack or complex of racks. UMIACS will provide the uplink for these TOR switches in terms of both physical and logical connectivity.

UMIACS research networks are currently based on a leaf/spine network infrastructure that currently supports 100Gb per second networking. UMIACS provide investment and support for the spine and PI(s) are responsible for making investments in the leaves. All leafs purchased are required to be redundant pairs and are often in port densities from 16 ports up to 64 ports per rack or rack complex. UMIACS will provide uplink for these leaf network switches in terms of both physical and logical connectivity.

Storage Nodes

Storage systems and their associated nodes that are purchased by PI(s) for their research must be purchased with a minimum 3-5 year initial hardware warranty. Storage nodes fall into two different categories: commercial closed source or open source storage systems.

At a minimum, UMIACS will provide data migration service for all non-scratch data on these systems as long as the previous system is under support and the target system is under support.

Commercial Closed Source

Commercial closed source storage systems must be under a software warranty and hardware warranty throughout their operational lifespan in UMIACS data centers. If renewal of the warranty period is declined, all data will need to be migrated before the expiration date of the warranty. After this expiration date, equipment will be depreciated and surplus through normal University procedures.

Open Source

For open source storage systems (Object Store, ZFS+NFS, XFS+NFS, etc.), nodes require a baseline 3+ year hardware warranty to be purchased. The UMIACS Technical Staff will provide best effort business hours support for nodes that are out of warranty and less than 7 years old. For systems with multiple components that are not co-termed in support contracts, staff assumes that all components are under warranty support. If one component is out of warranty, the whole system will be considered out of warranty.

For storage nodes that are greater than 7 years old and not under warranty, the UMIACS Technical Staff reserves the right to transition these nodes to store scratch data only, at any time after the 7 year mark, as defined by our Scratch Data Policy. This means there will be no further backups of this data, no single copy persistent data can be located on these nodes, and staff will make only a minimal effort to revive and repair the equipment. If more than 2 incidents are recorded within 30 days, or a major hardware failure occurs, staff reserves the right to notify the PI(s) that the node/system will be depreciated and surplus. The UMIACS Technical Staff will attempt to provide 30-60 days for migration or data retrieval. A best effort will be made, but the exact period staff is able to provide is highly dependent on the failure domain of the equipment.

Compute Nodes

Compute nodes are defined to be either contained within a scheduled resource environment (SLURM) or a standalone node without any UMIACS Technical Staff managed applications. Compute nodes must be purchased with a minimum warranty of 3 years. The UMIACS Technical Staff will support hardware throughout its usable life, as determined by staff, as long as the hardware is able to run a UMIACS Technical Staff supported operating system. Compute nodes will follow the active, extended, and legacy support categories.

The UMIACS Technical Staff will notify PI(s) within 30 days when equipment has its support contract expiring and ask if it should be renewed, if renewal is an option. If renewal is declined or unavailable, there is no action required for PI(s), and the equipment will move from active support to extended or legacy support, depending on age. For equipment that moves from the extended support category to the legacy support category, the UMIACS Technical Staff reserves the right to reorganize it within the compute infrastructure as needed. This ensures that staff can adequately support the transition of the workloads and potentially identify funding for new active supported nodes if appropriate.

PI(s) also have the choice to decommission nodes and remove them from UMIACS data centers at any support category. The UMIACS Technical Staff will need to be given 60 days notice to ensure that they can adequately prepare the node(s). If the equipment was purchased through UMIACS and is maintained on UMIACS inventory, the PI(s) will need to sign allocation form(s) for the equipment, reply to all requests for information about the location of the equipment for inventory purposes, and return the equipment to the UMIACS Technical Staff for disposal when it reaches its end of usable lifetime.