

Retirement Policy for Nodes

Policy Author: Jonathan Greenberg (jgrn@illinois.edu)

Passed: 6 Yes, 1 Abstain

Overview and Justification: Currently, purchased nodes on the Illinois Campus Cluster (ICC) are set for mandatory retirement after either 5 years (Taub) or 4 years (Golub). The first instance, Taub, is set to retire May 1, 2016 and the second instance, Golub, is set to retire September 3, 2017. There is currently no plan in place for what options are available to investors after retirement, aside from simply removing and returning nodes to the original investors, which has the possibility of significantly disrupting campus research and the ICC operations. This policy is to clarify the retirement process for Taub, Golub, and future investments in light of the nodes falling out-of-warranty, and eventually failing.

Technical Details:

Retirement Policy:

For all nodes purchased by an investor during an instance, once the node falls out of its guaranteed lifetime, defined as the warranty period of the node (for service periods, see *Overview and Justification* above), the investment enters into a new period, defined as the “retirement phase”.

An investor has two choices, at the point his/her investment enters the retirement phase:

- 1) The investor can elect to have the ICCP remove all or some of their node(s) and provide them to the investor. The specific node(s) returned to an investor will be matched (by serial number) to the specific node(s) that was purchased for that investor. If all of the investor’s nodes are removed, they will no longer have access to the primary or secondary queues.
- 2) The investor can elect to have the node remain part of the ICC, but agree to have their queuing access modified (see below).

If an Investor elects to remove their node(s), the Investor specifically acknowledges through written agreement that the ICCP is no longer responsible for the node(s), or for providing any support in relation to the node(s), once the node enters the investor’s possession. The ICCP will provide no guarantee as to the condition of the node once it is out of warranty. Investors removing all of their nodes will lose primary and secondary queue access to the cluster, although will still be able to access files in their storage space via the head nodes. The ICCP will be responsible for removing the investor’s node(s) from the cluster, at which point the investor will be responsible for picking up and relocating the node(s) from the Advanced Computation Building (ACB). The ICCP must remove the node(s) within one month of the request and contact the investor with instructions for pick up. For Taub, only the node(s) and all internal

components will be provided. For Golub, the node(s), all internal components, and the associated InfiniBand cable(s) will be provided. In both cases, equipment will be provided as “bare metal” without any packaging.

If the investor elects to keep their nodes that have entered the retirement phase as part of the ICCP, they must agree in writing to the following changes to their investor agreement:

- 1) An Investor’s access is guaranteed for an additional one year through September 1, 2017, with future extensions contingent on an Executive Committee vote.
- 2) Future renewals may require service fees, to be set by the Executive Committee.
- 3) An Investor will no longer be guaranteed a maximum of 4-hour waits to their investment, as over time the node failures will mean the total available nodes will be less than the original investment.
- 4) The maximum number of primary nodes an Investor/User can request will be equal to the number of originally purchased node(s) remaining as part the ICCP.
- 5) Primary queue will still allow the maximum walltime allowed by the ICCP.
- 6) The investor must maintain a minimum of 1 node as part of the ICCP to retain primary and secondary queue access.
- 7) An investor can, at their own expense, have their node(s) repaired or the warranty extended.
- 8) If an Investor elects to remove their remaining node(s) at a future date or if their remaining node(s) happens to be the node(s) that have failed, they will be returned the failed node(s), and will have their access to the secondary and primary queue removed.

Operators will maintain a list of node status linked to specific investors, so an investor can know how many functional node(s) still remain in their investment.

Fees: for the initial one year renewal, there will be no additional fees assessed for investors who elected to keep their nodes in the ICCP. Future renewals may assess fees at the discretion of the Executive Committee.

Final Retirement: This policy extends the initial timeline of the full retirement of a given node/instance by one year, after which the Executive Committee will vote to extend the retirement in one (1) year intervals beginning on May 1, 2016. The vote should take place at least 1 year before each renewal period is set to begin.

Proposed Timeline: May 1, 2016.

Impacts of the Policy:

Investors: after the policy is implemented, investors should decide if they want their nodes returned, or agree to the modified policy. If they remove their nodes, they will lose access to the primary and secondary queues. If they keep their nodes as part of the ICC, they will have a gradually reducing set of computational resources as the nodes fail.

Operators: the operators will need to maintain, as best as possible, ageing nodes. If an investor elects to have their nodes returned, the operators must provide the nodes to the investor within 1 month of the request. If an investor requests their node(s) be repaired, Operators will need to provide a quote and get billing information for the repair. Operators will need to retain a list of node status linked to specific investments.

Estimated Costs to Implement the Policy: 0.5 FTE annually of an ICCP Operator. This would cover in-place maintenance of the equipment, dealing with the increasing number of hardware failures/issues (particularly if an investor chooses to extend their warranty), and the eventual tear-down of the cluster and migration away from the DDN 10K disk system that is part of Taub.