Upgrade of the server architecture for the accelerator control system at the Heidelberg Ion Therapy Center

J. Mosthaf, K. Höppner, S. Hanke, S. Stumpf, A. Peters, T. Haberer (HIT)

Abstract

The proprietary accelerator control system of the Heidelberg Ion Therapy Center (HIT) runs on several classical server machines, including a main control server, a database server running Oracle, a device settings modeling server (DSM) and several gateway servers for auxiliary system control. As the load on some of the main systems, especially the database and DSM servers, has become very high in terms of CPU and I/O load, a change to a more up to date blade server enclosure with four redundant blades and a 10Gbit internal network architecture has been decided. Due to budgetary reasons, this enclosure will at first only replace the main control, database and DSM servers and consolidate some of the services now running on auxiliary servers. The internal configurable network will improve the communication between servers and database. As all blades in the enclosure are configured identically, one dedicated spare blade is used to provide redundancy in case of hardware failure. Additionally we plan to use virtualization software to further improve redundancy and consolidate the services running on gateways and to make dynamic load balancing available to account for different performance needs e.g. in commissioning or therapy use of the accelerator.

HIT Accelerator Facility

The Heidelberg Ion Therapy Centre (HIT) is a dedicated hadron accelerator facility for radio-therapeutical treatment of tumour patients. The two horizontally fixed treatment rooms as well as the experimental area can be served with proton and carbon beams with qualified beam parameters (MEFI), other ions like helium and oxygen have been tested.

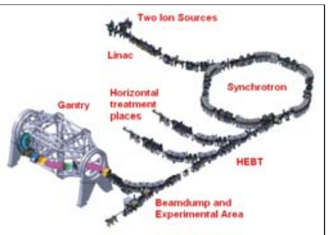

The achieved energy range of 88-430 MeV/u for carbon ions and 48-221 MeV/u for protons is sufficient to reach a penetration depth of 20-300 mm in water.

Parameter	Steps	Protons	Carbon
Energy	255	48-221 MeV/u	88-430 MeV/u
Focus	4(6)	8 – 20 mm	4 – 12 mm
Intensity	10 (15)	4·108-1·1010	1·107-4·108 1/s

Patient treatment in the two horizontal treatment rooms is running at approx. 40 patients a day and the experimental area is used in night shifts. The gantry commissioning is ongoing and expected to finish in early 2012.

Original Main ACS Servers

The original main ACS servers were housed inside two standard 19” racks together with secondary servers. We used Fujitsu-Siemens TX100S2 and TX200S3 servers with Xeon dual core processors and 2 GB of memory running Windows Server 2003 and Oracle 9. This was deemed sufficient for the ACS and ran from 2005 through the commissioning of the facility until 2010 when gantry commissioning began in earnest.

Underlying flaws in database design and installation as well as limitations in CPU power, RAM and I/O speed reduced performance of operations and calculations of device settings. One complete interpolation of all gantry devices took more than 4 hours per ion type. Downloading to the device control units (DCU) took more than **38 minutes** and flashing took another **30 minutes.** Cycle overhead (time between end of one beam cycle and start of a new one) was on average **1000 ms**.

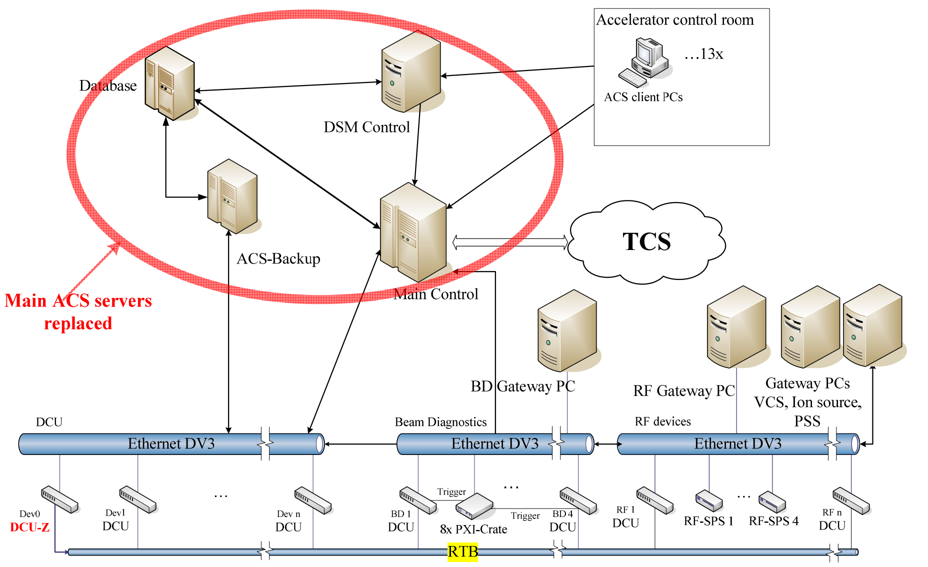

Server architecture and devices of the HIT accelerator control system with replaced servers shown.

New Blade Center

Replacing the ageing server structure with a new, state of the art blade center. To spread the costs over several budgets, it was decided to first only replace the main ACS servers more or less directly with blade servers and also procure the necessary infrastructure for further expansion. We started with one **16-space enclosure** with redundant power supplies, management modules and dual **Flex-10 10Gb/s** Ethernet connections.

	Original System	New Blade System
CPU/Cores	1x2 cores	2x8 cores
CPU speed	1.6 GHz	2.53 GHz
RAM	2 GB	24 GB
Operating System	Windows Server 2003/ Windows XP 32bit	Windows Server 2008 64bit
Network connection	1Gb/s	1Gb/s ext. 10Gb/s int.

Four identical blade servers with **two 8-core CPUs, 24 GB of RAM** and two internal HDDs and one **storage blade** with four HDDs in a RAID configuration are integrated into the enclosure. The new database server, running Oracle 11 on Windows Server 2008, is connected to the storage blade and has a completely new configuration. All blades are interconnected and share the same power supply and network.

The enclosure and all built-in blades can be controlled and configured via web server and remote desktop connections. New blades can easily be connected and integrated into the system. For future expansions, we still have 11 slots free in the enclosure.

Performance improvements

The table to the right shows the rough improvements in DSM calculations and cycle overhead with the old servers and the new blade center. The improvements are apparent in all figures.

	Original system	New blade system	Improved by
Interpolation incl. gantry devices	~4-6 hrs	~1.5 hrs	>200%
Download per ion type	~40 min	~10 min	>400%
Flash per ion type	~30 min	~8 min	~275%
Avg. Cycle Overhead	~1000ms	~740ms	~35%

The scatter graphs below show changes in cycle length and cycle overhead. The red green data shows cycles after the upgrade. Especially the spread of the overhead time is superior.

Cycle length

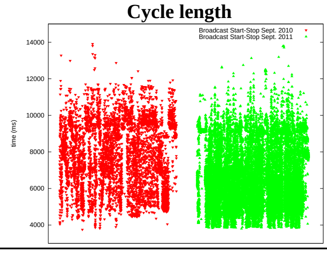

Cycle Overhead

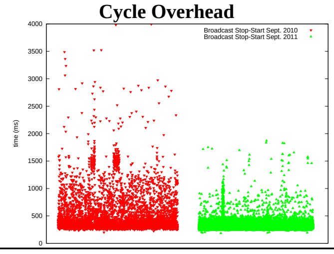

Conclusion and Outlook

The performance gains with the new blade servers are substantial. We show significant improvements in DSM computations and in operating. Waiting times for database queries are vastly reduced. Redundancy is enhanced by the hot spare blade server which is pre-configured to run all ACS services and by the dual Ethernet connection.

For the future, more blade servers are planned to increase redundancy and to enable virtualization and load balancing of servers that would profit from it (DSM, gateways…)