Appendix B: High Energy Physics Scenario generated for the Workshop on New Visions for Large Scale Networks[1]

Never before has the scientific mission of particle physics research been so dependent on state-of-the-art information technology. Collaborations of hundreds to thousands of physicists and engineers are formed to create accelerators, detectors, and analysis systems with a productive life of tens of years. These analysis systems form a complex and widely distributed "fabric" of computing and storage resources.

The non-deterministic nature of quantum physics, uneasily understood during the last century, inevitably requires the measurement and analysis of billions of particle interactions to observe and understand fundamental processes. Particle physics experiments have pushed against the limits of technology, electronics, computing, and networking for decades. Detectors with millions of channels, each recording precise amplitudes with a resolution of picoseconds, have in the course of 40 years succeeded detectors with a few single-bit "yes/no" measuring devices. Information flows from such a detector at up to a terabit per second and must be drastically filtered in real time because of limited storage, analysis, and networking facilities.

The Large Hadron Collider (LHC) experiments at the European Organization for Nuclear Research (CERN) will rapidly reach tens of petabytes of stored data under intense analysis. The design, construction, and data analysis for an experiment require the combined intellect and dedicated work of international collaborations. However, technological limitations on the storage, transmission, and analysis of data impose difficult, even dangerous choices. For example, the LHC experiments expect to be able to record and share over networks less than one millionth of the collisions they observe. This draconian real-time selection will necessarily have to be optimized for "somewhat expected" new discoveries rather than the "totally unexpected" ones that are the dream of every scientist.

Even after the draconian selection, LHC collaborations will face the challenge of empowering thousands of geographically distributed physicists to use their intellect and wisdom to derive physics insight from tens of petabytes of data. Although the raw cost of bandwidth is no longer a crippling impediment, the end-to-end performance of applications is often unacceptable. Success will rely on middleware research to support data-intensive, worldwide collaborative science that is only beginning. A minimum requirement is the location-independent ability to analyze data to empower all of an experiment's physicists to work collaboratively on databases, growing to tens of petabytes in 2005-2010, using all computing resources to which they have access.

However a qualitative change in the way research is performed would be enabled if we could free the real-time selection of data from the constraint of a single filter system with

---

[1] "Workshop on New Visions for Large-Scale Networks: Research and Applications", Large Scale Networking (LSN) Coordinating Group Of the Interagency Working Group (IWG) for Information Technology Research and Development (IT R&D), March 12-14, 2001, Vienna, Virginia

selection algorithms decided by committees. The availability of networks with end-to-end terabit performance could make this possible, but speed alone is not enough.

The data acquisition and filtering systems might profitably become geographically distributed and operate as highly parallel, largely asynchronous data flow systems. The terabit systems that will become operational in 2005 for the LHC will include a multi-terabit capacity switching fabric, but individual data acquisition nodes and filtering nodes will communicate at gigabit speeds. In addition to the substantial bandwidth requirement, challenges include:

- The multicast service required when more than one remote filtering center is available.
- Achieving adequate error rate and robustness without *ever* allowing the implementation of the "wild idea" to impact the detector-site data acquisition system

Based on this scenario, the participants in the workshop generated the requirements given in Appendix C.