An Overview of Partners HealthCare Clinical Information System Architecture

by Ethan Fener

History

In the mid-1980's a MIIS-based, traditional HIS was implemented at the Brigham & Women's Hospital (BWH), which employed character cell terminals and serial connections to mini-computers. To overcome the maximum connection limitations of this platform, a new client/server PC-based architecture was developed from 1989 to 1991. It was known as "the new platform" or TNP [see figure 1]. This second generation information technology (IT) architecture for the hospital accommodated the explosive demand for data, access and applications. It continues today servicing well over 4000 desktop PCs as clients.

In 1993, BWH and Massachusetts General Hospital (MGH) joined to form Partners HealthCare System in response to the business changes of the healthcare industry. These two hospitals were the anchors with more hospitals and ambulatory care centers joining over time. The vision set forth was an integrated network of thousands of physicians delivering care to perhaps 2,000,000 covered lives. When speaking publicly about this new phenomenon of integrated healthcare delivery networks, the late Richard Nesson, M.D., who served as the first president of Partners, stated that "what will separate the good from the great are world-class information systems."[1] From an IT standpoint, the estimate was that in 2000, some 50,000 desktops would be tightly or loosely coupled to the network with 20%, or 10,000 desktops, accessing clinical information at any time.

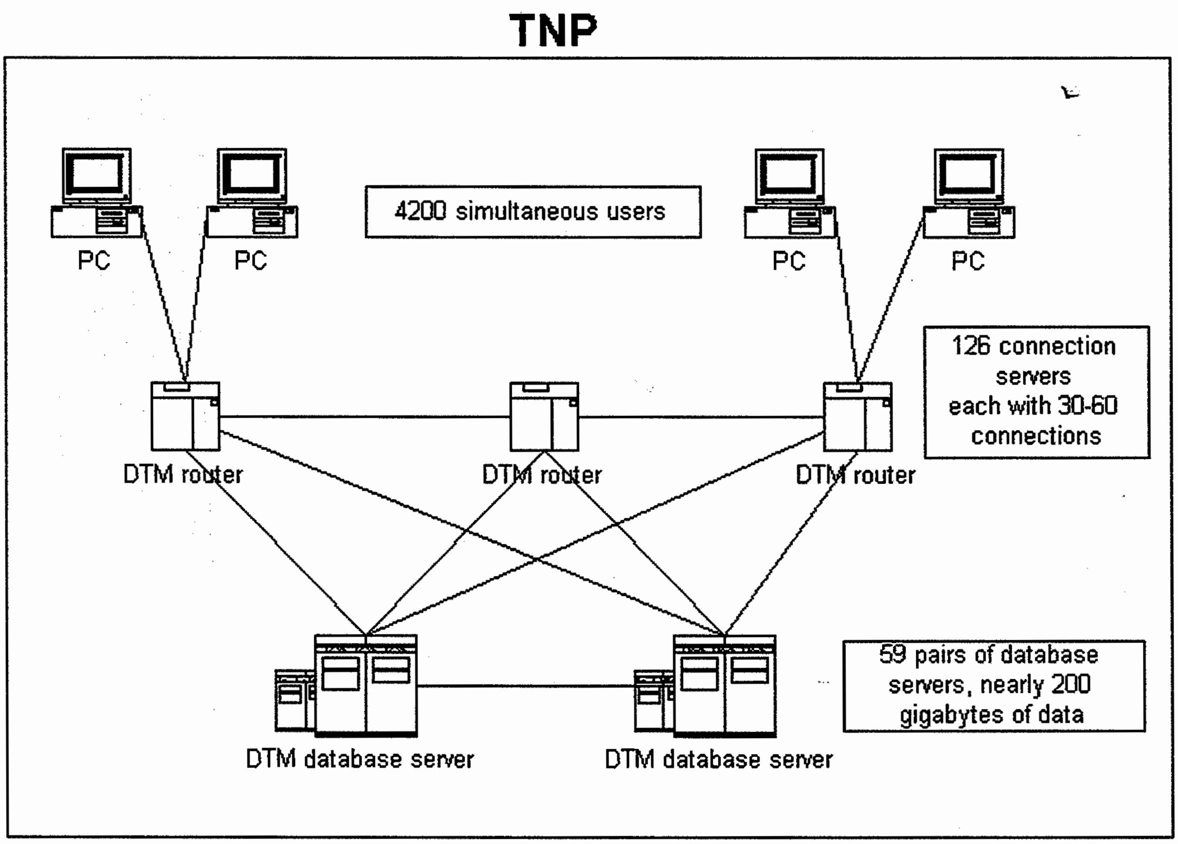

Figure 1

Architecture

P3 was the new information system architecture conceived in the mid-1990s to meet this challenge [see figure 2]. The highly-acclaimed three-tiered approach at BWH[2,3] was enhanced by emphasizing industry standards, increasing reliability and by choosing more extensible underlying hardware and software. NT was selected as the operating system for servers and the client network, i.e., user/device authentication, e-mail, file sharing and print queue management. Visual Basic (VB) and Microsoft's component object model (COM) were selected for new application development. InterSystems offered a migration strategy from DataTree MUMPS (DTM) to Caché (previously openM) that preserved the investment in existing code, programmer training, database design, and the success of earlier efforts. This included the VISM.OCX which exposed the M database to the VB environment.

A design limitation of TNP was the permanence of the connection between the client and the middle tier DTM server/router. If a server/router failed, the 30 to 70 attached clients were forced to stop and start their DTM client processes, which meant losing the last transaction. M Services was the new middleware written into P3 to arbitrate client requests for data and make the application servers fault tolerant [see figure 3]. VB applications first call MSERVICE.DLL. Thereafter, M services controls the VISM.OCX on the client and establishes and manages communications to the least busy application server to access Caché. In this design, clients retain context and the middle tier is stateless. Should an application server fail or need to be taken out of service for maintenance, the load-balancing algorithms within M Services will re-direct client communications to another application server on the next request for data, making device failures imperceptible to interactive users.

Since clients maintain state information, an application server is an arbitrary destination. Therefore, M Services must maintain a variety of usage statistics on active processes for better resource management. Currently, it monitors the connections from clients to application servers and destroys them after 30 minutes of inactivity. The number of open connections to an application server at any time is therefore a reasonably accurate (i.e., slightly overstated) measure of simultaneous users.

Partners is not immune to DLL conflicts and incompatibilities that occur when Microsoft Office applications,

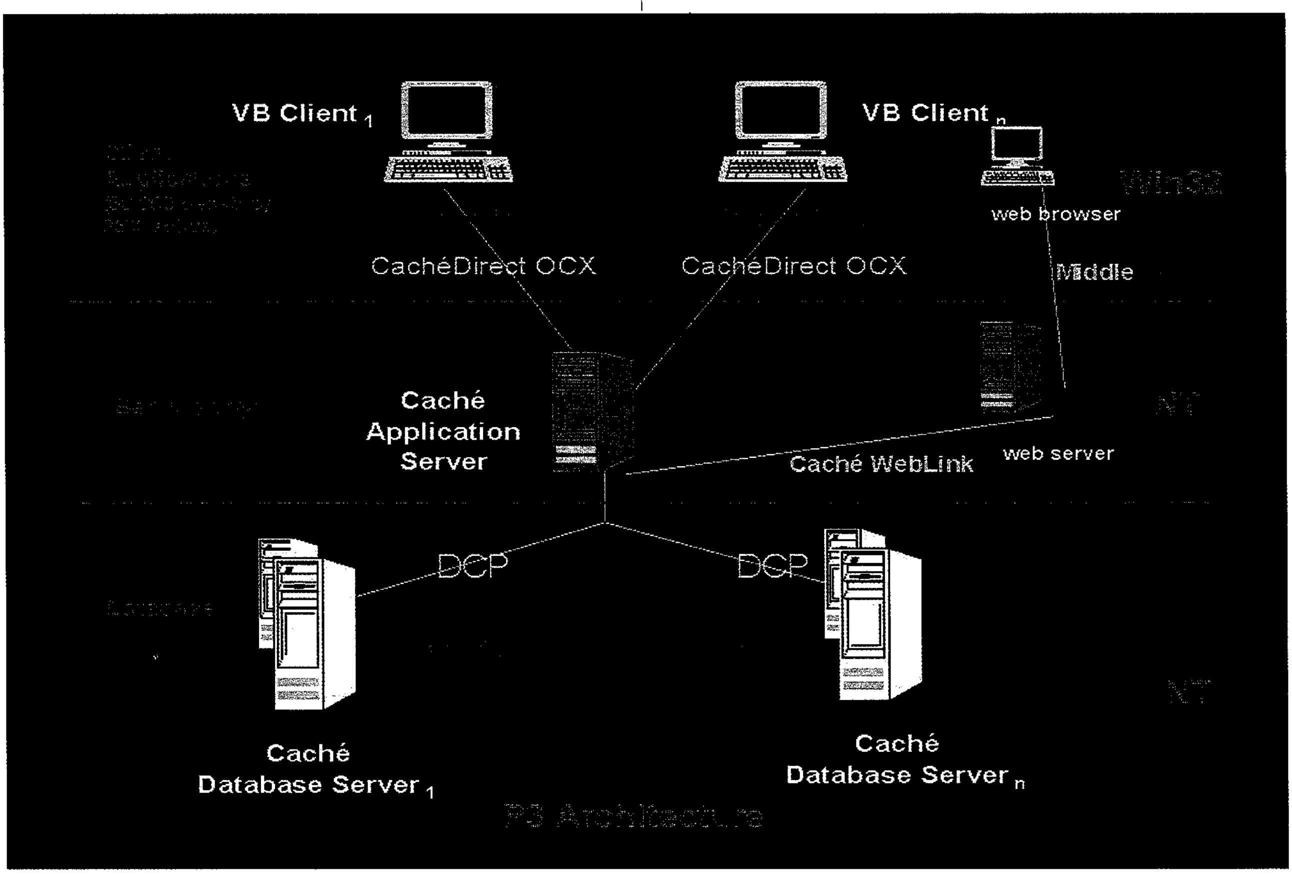

Figure 2

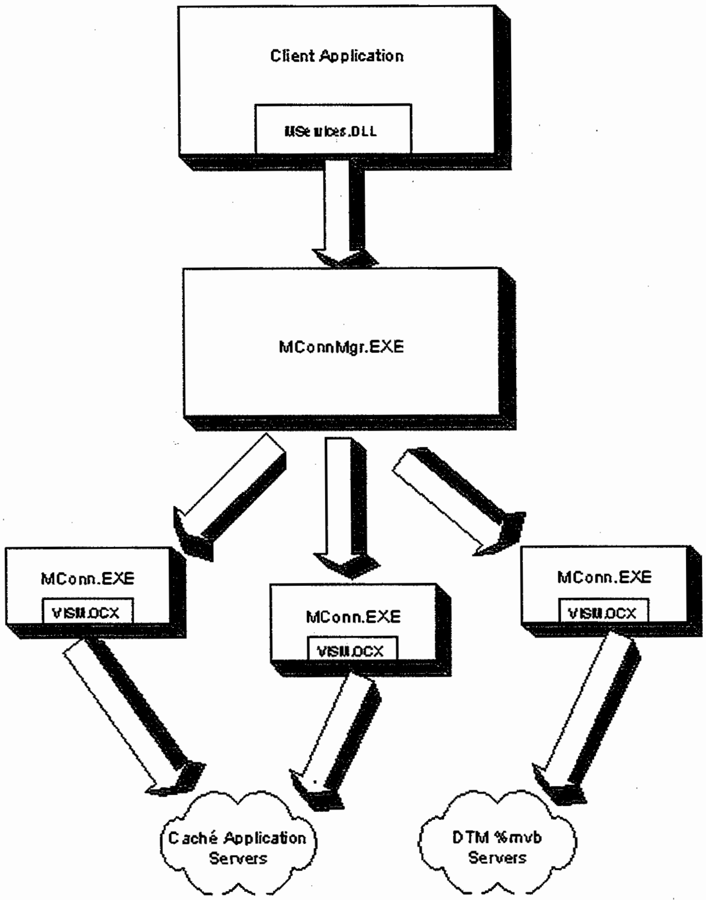

Figure 3

Windows 95, internally developed applications and purchased third party applications come together on the same desktop. Most of these problems are kept in check because of the Partners customized desktop and a sophisticated management strategy. A few clinical applications however, have been browser-enabled to accommodate the loosely-affiliated clients, which are owned and managed by other entities. These requests for data are serviced by the Weblink add-on to Caché coming through standard HTTPD web servers.

The application servers protect the database servers. All interactive user processes and batch jobs execute only on the application servers. Global data move via distributed Caché protocol (DCP) over UDP between the databases and the application servers and via TCP/IP to the client. The database is redundantly stored on a shadow server, a hardware twin of the live system, by using a native Caché shadow process.

Typical hardware configurations for database servers are: two 400 MHz Pentium II Xeon CPU's; 512 megabytes of memory; 41 or 82 gigabytes of usable RAID 5 storage; 4-8 gigabytes RAID 1 storage for NT 4.0 enterprise; 8 gigabytes RAID 1 for journal space; a 100 megabit switched Ethernet network connection and NIC; a 16 megabit token ring connection (to DTM) and two DLT tape drives (shadow server only).

Users see minimal downtime on the live database servers. Daily backup is taken from the shadow server first by stopping the shadow process that is sending transactions from the live server, then by stopping Caché to close the database. All disks are backed up to tape via ARCserve, a widely used third-party backup product. After backup, the shadow server returns to its normal dejournalling state. A background process (SHADCOMP) performs a node-by-node comparison of all global data on each system. Any discrepancies detected are recorded and reported via e-mail. SHADCOMP can reliably check 20-25 gigabytes of data per day in both directions without negatively impacting user performance. Database integrity checks are automatically run weekly on all Caché systems and reported via e-mail.

Should a live machine fail or require periodic maintenance, failover to the shadow is initiated. Very simply, the first step is to confirm that the last SHADCOMP completed successfully and that shadowing from the live server is up-to-date. Caché is then stopped gracefully on the live server and NT is shut down. The shadow server is removed from the domain and its network identity (node name and IP address) is changed. It is rebooted as the live server and rejoins the domain. The failover process has been carefully refined and routinely takes 12-15 minutes with some manual intervention.

DTM Migration

TNP serviced BWH well for several years, but by 1997 it reached the limits of its DOS underpinnings with respect to expansion, data storage and disk I/O. It was architected for 3000 clients and could not run long-term at 130% of capacity. At the same time, Y2K research and remediation revealed a tremendous capital asset of more than 45,000 DTM routines. Preserving this code base mandated that the upper two tiers remain unchanged [see figure 4]. The solution then was to migrate the existing 180 gigabyte DTM database to P3 with minimal impact to the 4200 users. Such a project demanded rigorous planning and testing to address issues such as:

- What is the I/O profile of the existing datasets?
- How much data should be stored on a single server?
- What is the network thru-put of Caché in DTM compatibility mode?
- How do we insure the integrity of the data through the transition?

DTM Migration

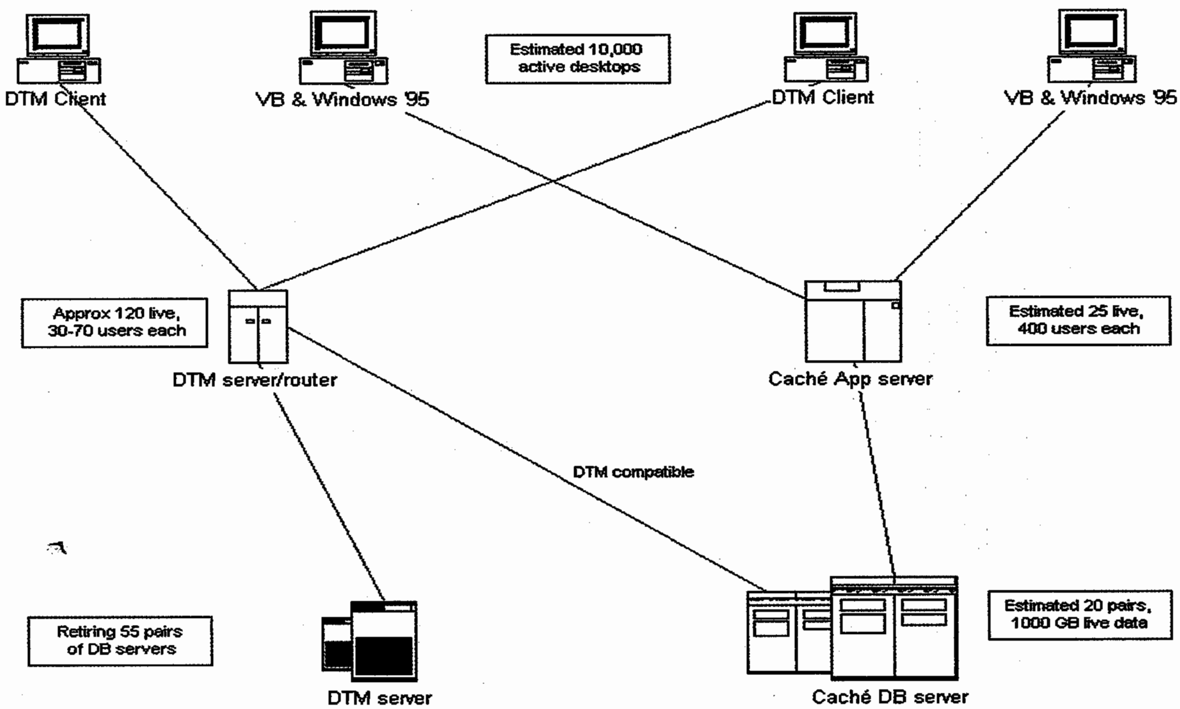

Figure 4

- Do DTM applications behave differently when data requests are serviced by Caché?
- Is there a failsafe should the unexpected occur and we need to go back to the DTM server?
- What will it take in time, labor and capital to migrate all the data?

Over the better part of a year, these questions and others were answered and will be discussed more fully in a future writing. The migration commenced and approximately 52 gigabytes of data have been migrated at this time. The project is slated for completion in calendar 1999.

Usage

A variety of institution-specific and enterprise-wide VB applications are flourishing on P3. An abbreviated list includes:

- Clinical data repository (CDR). An ever-expanding ingest of patient data from BWH and MGH. Feeds from other affiliated hospitals are planned.
- Order Entry (MGH OE). A VB re-write of the successful physician order entry application deployed at BWH tailored to meet MGH needs.
- Longitudinal Medical Record (LMR). A results viewer that displays patient data from the CDR for the ambulatory care setting.
- Enterprise-wide Master Person Index (EMPI). A demographic database that other applications use to look up patients by institutional medical record numbers.
- Partners Paging Directory (PPD). A browser-based application for looking up employee information that includes online paging. This is especially useful for people who work on both hospital campuses.

As of this writing, three Caché application servers handle about 750 client users with a negligible number of Telnet sessions. One server is reserved for background jobs and receives a small fraction of interactive users. Approximately 500 gigabytes of global data are stored on the live Caché servers and duplicated by the shadows. Users and storage are expected to double in the next twelve months as the existing applications expand into more areas of the Partners affiliated institutions.

Conclusions

NT server 4.0 lacks the robustness and maturity of operat-