

# Clusters to pay their way

**F**irst there was the LAN, then the WAN, now it's the SAN – the system area network. Originally developed by Tandem for its UNIX-based Himalaya servers, the new architecture is at the heart of its recently released enterprise-class NT servers. They are part of Tandem's strategy to bring its experience as a fault tolerant transaction processing specialist to bear on the brave new world of Internet commerce. On-line transaction processing (OLTP) now becomes, in the company's phrase, 'Internet transaction processing' (ITP). Having already captured most of the OLTP market – Tandem claims its technology is behind 80 per cent of the world's ATMs (automated teller machines), 66 per cent of credit card transactions and 90 per cent of the world's securities transactions – the company believes it is time to take money from the Net. Its new range of S-series servers is designed to ensure secure financial transactions over the Internet and includes three new Windows NT-based servers.

The Pentium Pro 200MHz-based servers each run Windows NT 3.51 or 4.0, include up to 1Gb of memory and on-board Redundant Arrays of inexpensive disk controllers enabling on-line expansion of disk files. They can also come equipped with the optional CAS (Cluster Availability Solution) – cornerstone of Microsoft's Wolfpack NT clustering initiative. With the CAS failover technology, two WNT Server nodes can reliably and transparently recover from hardware, application or operating system failures. The cluster-enabled WNT servers include the two processor S100 or four processor S1000 and S1000RM(rack-mounted). In addition, Tandem is porting its suite of middleware and database applications ServerWare – including NonStop SQL database and transaction processing monitors – to the WNT Server-based systems in 1997. Through its ServerNet technology, Tandem believes companies can build low-cost, reliable computer clusters that can grow in terms of processing, storage or throughput. The inherent difficulty of most clustering approaches is that they offer limited scalability. Traditionally, the problem has been that as extra elements are added more of the system resources – largely processor resources – are used in managing the cluster itself until each additional element costs more than it adds.

A SAN architecture avoids processor overload by minimising the use of the processor for transactions other than processing. To support the coming generation of data intensive, media rich (audio, video, graphics etc.) apps it's necessary to not only process data quickly but also move it efficiently through the system. "The challenge of deploying tomorrow's most powerful applications is no longer processing

speed, availability, storage capacity or cost. The new challenge is server throughput, i.e. efficiently moving extraordinary amounts of data within the server," says Kurt Friedrich, general manager of Tandem's Parallel Systems group. ServerNet technology does that by not forcing data through the processor at every datapath. Unlike bus-based systems, SANs cluster all major components of the system as independent elements. The result is that any element – processor, disk or I/O – can interact with any other element without processor intervention. It means, says Tandem, faster response times, freed up processor cycles, unparalleled system and I/O bandwidth and inherent reliability. ServerNet has already been taken up – or is being evaluated – by 85 per cent of PC-server manufacturers, including Compaq, NEC and Dell. Tandem has also recently announced a new collaboration with Siemens Nixdorf and its subsidiary, Pyramid Technology, to work towards an increasingly open standard for cluster technology for both the UNIX and NT server environment.

ServerNet drivers are also designed to be fully compatible with Microsoft's forthcoming WNT clustering software, Wolfpack, when it debuts early in 1997. In fact, phase one of the Wolfpack project only allows for two node clusters. The beta of phase 2, supporting larger clusters, providing higher scalability and additional load-balancing capabilities is only due for release in 1998.

And there's the rub. For truly enterprise fault tolerant computing, Tandem are relying on ServerNet linking the NT-based servers with the UNIX-based Himalaya servers. It's almost UNIX by the back door. While the WNT-based S series servers have only four Pentium Pro processors, the flagship Himalaya servers (S7000 and S70000) can scale from two to more than 4,000 processors. By using the same ServerNet technology across the UNIX-based and NT models, Tandem claims it offers the best of both worlds. For Himalaya users it means their critical applications can take advantage of the WNT server platform's cost-effectiveness, breadth of hardware, ease of development and ease of use. For WNT users it means their most critical, high volume, data intensive apps will have access to the Himalaya's massively parallel processing power, scalability, throughput, and continuous availability. The Himalaya servers are seen as the backend powerhouse for clusters of front-end WNT server based systems – such as in a bank where NT is used for the branch office front-end while customer account information is kept on a central Himalaya. But for the moment it is NT that goes to the mountain rather than the reverse – full-scale enterprise clustering is one peak that NT has yet to scale.

– Roger Laing

