VMAX 40K Description physical:

Symmetrix VMAX 40K dispersed system bays allow a single VMAX 40K to be separated across two locations by up to 82 feet (25 meters) enabling deployment in dense data center environments where floor loading and other physical limitations and obstructions would otherwise limit configurations.

The Symmetrix VMAX 40K system architecture provides a Quad Virtual Matrix™ with the core element being the Symmetrix VMAX 40K engine. The VMAX 40K Engine includes up to 256 GB cache memory per Engine, front-end connectivity, and backend connectivity.

The engines and standard or high density disk array enclosures are installed in a specific order, beginning with engine 4. The order of engine population is the same for both standard and dispersed layouts. Additional disk array enclosures can be daisy chained into 30 and 45-drive loops for standard configurations, 50-drive loops on high density configurations, and 65 drive loops on mixed configurations

In all cases, engines 4 and 5 are added first, followed by engine 3, engine 6, engine 2, engine 7, engine 1 and engine 8. Engines added to dispersed configurations are added in the same order as those in a standard configuration, but may be placed in a second system bay.

Storage bays and DAEs are populated in the same order as the engines to which they connect. Direct connect bays are added before daisy-chain bays. However, because high density DAE storage bays do not support third-level daisy chains, configurations with high density DAE storage bays do not include bays 3C and 3D. This is also the case when mixing standard and high density drive bays within a system.

System capacities

Symmetrix systems are expanded by adding engines, disk array enclosures, and physical

memory. Table 2 on page 19 shows system capacities.