Table of Contents

Super Micro Computer, Inc.
980 Rock Avenue
San Jose, CA 95131 USA
www.supermicro.com

Reference Architecture

Supermicro 2U 4-Node TwinPro² Reference Architecture for EsgynDB®

Delivering on Converged Big Data Platform

Executive Summary

Supermicro and Esgyn have teamed up to bring to market a Big Data Appliance that delivers faster business insights using EsgynDB, a Big Data converged database optimized to run on the Supermicro 2U 4-node TwinPro2 SuperServer platform. This Big Data Appliance comes complete with the software stack pre-installed, pre-configured, and pre-tuned for the use cases of Modernization of Data Lakes and Acceleration of IoT value, and to deliver the highest performance on standard benchmarks like YCSB, TPC-C and TPC-DS. In fact, all 99 queries of the TPC-DS benchmark run straight out of box, and EsgynDB is the only product on Hadoop that can do this. EsgynDB is tightly integrated with Hadoop and supports HDFS, HBase, ORC, and Parquet. The Appliance is complete as-is or can be leveraged as a building block and scaled up to any cluster size needed.

Adoption of the Supermicro-Esgyn Big Data Appliance significantly reduces the complexity in deploying a Big Data solution to modernize data lakes, to support hot and cold data IoT analytics, and to create a converged database platform for future needs that include support of multi-varied data and OLTP, ODS, BI and analytical workloads.

EsgynDB, with its mature ANSI SQL interface on Hadoop, running on the SuperServer platform, can be used as a vehicle not only to rapidly execute pilot or proof of concept projects, but also to clearly demonstrate acceleration of the ROI from Big Data investments.

Target Audience: This document is intended for decision makers, system and solution architects, system administrators and experienced users that are interested in reducing design time or simplifying the purchase of a big data architecture containing both Supermicro and Esgyn components. An intermediate knowledge of Apache™ Hadoop® and scale-out infrastructure is recommended.

Document Purpose: The purpose of this document is to describe the appliance bundle of EsgynDB Enterprise on Supermicro cluster.

Why Supermicro-EsgynDB Appliance

The Big Data Appliance comes complete with the software stack pre-installed, pre-configured, and pre-tuned, from the OS, to Hadoop, to EsgynDB, all tested, and certified by Esgyn Labs, on a complete Supermicro hardware platform that includes servers, memory, storage, and network adapters, all chosen for optimum performance and reliability, targeting the use cases of Modernization of Data Lakes and Accelerating IoT vale.

It is sold as a building block unit with built-in redundancy for high availability and positioned as a single converged platform for a spectrum of DBMS workloads from online transaction processing (OLTP) to operational data store (ODS) to data warehousing (EDW) to business intelligence (BI) to online analytical processing (OLAP).

Available on free trial basis for up to 45 days.

Introduction

This reference architecture has been created to assist in the rapid design and deployment of EsgynDB Enterprise software on Supermicro infrastructure. It is also intended to identify the software and hardware components required in a solution to simplify the procurement process. The recommended Supermicro appliance has been carefully tested with a variety of I/O, CPU, network, and memory configuration for transactional and analytical workloads.

Solution Overview

Supermicro and Esgyn solutions enable people to derive new business insights from Big Data by providing a platform to store, manage and process data at scale. This reference architecture provides several performance-optimized configurations for deploying EsgynDB Enterprise clusters on Supermicro infrastructure. When compared to previous generations, the new configurations provide a significant reduction in complexity and a recognized increase in value and performance.

Configurations include preinstalled, pre-configured, and tested versions of Linux, Hadoop, and EsgynDB on the Supermicro server platform – a highly optimized server with the right mix of compute and storage performance characteristics for parallel processing systems, like Hadoop. The configurations selected are the result of a great deal of testing and optimization done by Esgyn engineers resulting in the right set of software, drivers, firmware and hardware to yield excellent workload performance characteristics.

Supermicro/Esgyn Big Data solutions provide excellent performance and high availability, with integrated software, services, infrastructure, and management – all delivered as one tested configuration to significantly reduce deployment overheads and operating expenses.

Supermicro TwinPro Server Architecture

The Supermicro TwinPro architecture is based on the Supermicro proven Twin technology to provide exceptional throughput, storage, networking, I/O, memory and processing capabilities in a 2U form factor. Customers can further optimize Supermicro solutions to resolve the most challenging IT requirements and benefit from exceptional Total Cost of Ownership (TCO).

TwinPro solutions are designed for simplified deployment and maintenance, and assembled with the highest quality to ensure continuous operation even at maximum capacity. Customers in high-end enterprise, data center, HPC and Cloud Computing environments receive the greatest competitive advantage from data center resources with the Supermicro TwinPro.

- **Performance**: The latest generation of Intel Xeon processors significantly heighten performance of Supermicro TwinPro solutions. With up to 28 cores per socket, 6 channels of DDR4 and up to 2TB of DDR4-2666 in 16 DIMM slots, memory speeds will reach a maximum of 2666MHz. Supermicro X11 and X10 TwinPro systems are optimized to deliver incredible performance to high performance environments and will handle complex workloads with ease.
- **Flexibility**: High flexibility and various networking options are key features of

the TwinPro. Each node includes a Supermicro I/O Module (SIOM) for networking options such as 1G, 10G, 25G, 40G or 100G and 2 low-profile PCI-E 3.0 x16 expansion slots, giving our customers the flexibility to include add-on cards when they are required.

- **Efficiency**: The Supermicro Twin Architecture product line delivers the industry's densest, cost -efficient servers. The TwinPro doubles the performance of a traditional 2U system and houses four nodes to reduce data center footprint and includes redundant Titanium Level Digital (96%) Power Supplies for maximum power efficiency.

The reference architecture leverages the Supermicro X10 generation TwinPro^2 platform. It hosts 4 independent hot-swap 2-socket systems (nodes) in a 2U chassis. The system is powered by redundant 2000W Titanium Level (96%+) power supplies.

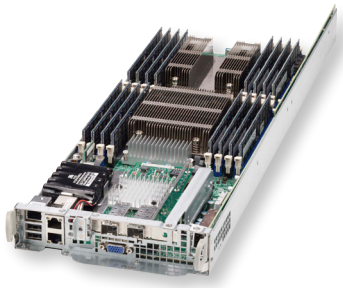

Node view (1 of 4)

EsgynDB Enterprise Overview

EsgynDB offers a unique value proposition by supporting the full spectrum of database workloads at Big Data scale on a single converged database platform, simplifying your database landscape, reducing or eliminating data movement, and reducing overall latency. EsgynDB supported workloads include:

- Online transaction processing (OLTP)
- Operational Data Store (ODS)
- Data Warehousing (EDW)
- Business Intelligence (BI)
- Online Analytical Processing (OLAP)

EsgynDB reduces ETL costs by transforming data (ELT) within the database. It's Massively Parallel Processing (MPP) architecture ensures that the strictest SLAs can be accomplished by executing queries in parallel. A mature ANSI SQL engine enables application portability; proven to handle petabytes of data at high concurrency in numerous demanding enterprises for over 20+ years.

Here is an overview of EsgynDB capabilities:

- Comprehensive and full-functioned SQL support allows companies to reuse and leverage existing SQL skills to improve developer productivity
- Easy integration with other ecosystem components to extend the data sources available from streaming sub-systems to sentiment analysis and so on
- Support for hot and cold data
- Join structured, semi-structured, and unstructured data from disparate data sources all within one SQL query
- Interoperability with new or existing applications and 3rd party tools via support for standard ODBC and JDBC
- Provides full ACID transactional protection across multiple rows, tables, and SQL statements
- Sophisticated and massively parallel database engine supports complex queries at high concurrency and throughput

- World class Optimizer with abilities to tune the degree of parallelism to suit individual queries and the scale or cardinality of data being queried
- Supports full active-active distributed transactions across data centers to scale workloads with zero lost transactions for disaster recovery (EsgynDB Advanced edition feature)
- Supports multi-tenancy to best optimize the available resources to meet SLAs of different tenants (EsgynDB Advanced edition feature)

This breadth of capability, referred to by Gartner as Hybrid Transactional and Analytical Processing (HTAP) and by Forrester as Translytics, makes EsgynDB ideally suited to a range of use cases, including:

- Internet of Things (IoT)
- Data Lakes

Big Data Workloads

Typical Hadoop workloads can be broadly categorized into 4 different types: Batch, Non-Interactive, Interactive, and Operational. These categories vary greatly in terms of their response time expectations, concurrency, as well as the amount of data that is processed. Hadoop is well positioned to address BI and analytical workloads. Supporting Operational workloads on Hadoop is an emerging Hadoop market category and Esgyn is the only company that can do this. Traditionally these workloads have been relegated to the domain of traditional relational databases. But there is growing interest and pressure to embrace these workloads on Hadoop due to Hadoop's benefits of significantly reduced costs, reduced vendor lock-in, and its ability to seamlessly scale to larger workloads and data sets.

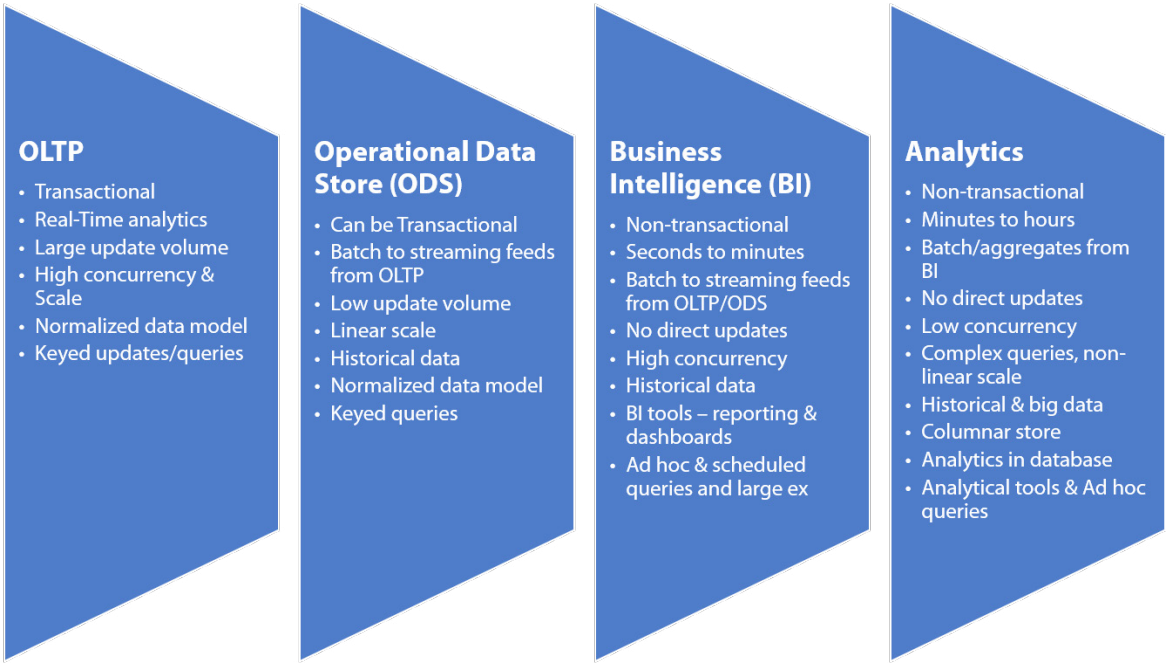

Internet of Things (IoT)

Most IoT systems deal with real-time or near real-time data. Analysis of real-time data together with older data on a single platform enables well-informed responses. In technical terms, this could be construed as combining hot data with warm, and even cold data, on a single platform. If this can be accomplished with minimal data movement and data transformation, then the overall response times could be significantly reduced. This results in a system that allows not only massive ingestion but also real-time reporting at the same time, and at significantly lower cost. The platform needs to deliver many critical aspects such as scalability, availability, and high concurrence. EsgynDB can do exactly that, maintain sustained ingestion of transaction protected incoming data while concurrently reporting or alerting on that same data. No batch updates. No data latency. The Supermicro-Esgyn Big Data Appliance makes scaling out an IoT solution simple.

Data Lakes

Pumping any and all data into data lakes has not worked. Most of the data lakes are dry or at best ponds or swamps as Business Units are not willing to contribute their data. This is because they are not able to see the same business reports with SLAs they are used to. Currently their world is used to SQL interfaces for pre-canned or ad hoc queries, and self-served Business Intelligence reports. Even after allocating significant resources (Java programmers, data scientists, Hadoop experts, or people with vendor-specific skills) data lakes are not delivering to the expectations of business users. The promise of Hadoop-based data lakes may have been oversold and over purchased. Enterprises are pondering issues such as: whether a data lake itself has become a silo instead of breaking down silos; whether data lakes are moving applications away from data; or whether schema flexibility is really helping to make life easier for business users. In almost all implementations, lack of a mature SQL engine is the common denominator why enterprises are not able to see the promised return on investment. Adoption of the Supermicro-Esgyn Big Data Appliance significantly reduces the complexity in deploying a Big Data solution to modernize data lakes. Get real value and actionable insights from your data lake by bringing the workloads to the data lake. Most other database engines on Hadoop lack the maturity and sophistication to serve the full spectrum of workloads needed to maximize the potential of your data lake. From operational to analytical, EsgynDB is the one database engine you need. The Supermicro-Esgyn Big Data Appliance makes building a data lake easier and maximizes the return on that investment.

EsgynDB innovations built upon Big Data software stack

EsgynDB is designed to build upon and leverage Apache Hadoop and HBase core modules. Operational applications using EsgynDB transparently gain Hadoop's advantages of affordable performance, scalability, elasticity, availability, etc. Figure 1 depicts a subset of the Hadoop software stack. Items in green are specifically leveraged by EsgynDB, such as HBase and HDFS. To this stack, EsgynDB adds (items in orange) ODBC/JDBC drivers, the EsgynDB database software, and a distributed transaction management (DTM) subsystem for transaction protection.

EsgynDB delivers innovation on top of Hadoop in these key areas:

- A full-featured ANSI SQL implementation whose database services are accessible via

a standard ODBC/JDBC connection

- A SQL relational schema abstraction
- Distributed ACID transaction protection
- Full cross data-center active-active distributed transaction support, to scale reads and writes, support local access and comply with safe harbor rules, with zero lost transactions in a disaster
- Performant response times for transactions comprised of both reads and writes
- Parallel optimizations for both transactional and operational reporting workloads

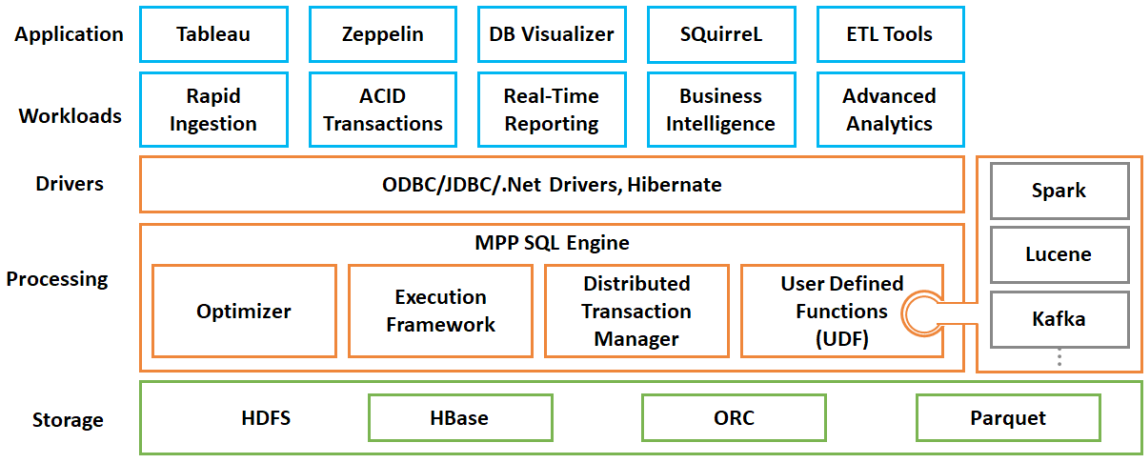

Figure 1. *EsgynDB Ecosystem*

Solution Components

Cluster Configuration

A number of key factors were considered prior to designing EsgynDB Big Data appliance. The following subsections articulate the design decisions in creating the baseline configurations for the reference architectures.

Operating System

Red Hat® Enterprise Linux®

Computation

When sizing the processing power for an EsgynDB Enterprise Supermicro cluster, concurrency and types of workloads are the other key considerations for number of nodes. The number of nodes and cores reflects the amount of parallelism available for concurrent users of the applications running on the cluster. If typical workloads are high-concurrency short queries, then nodes with fewer number of cores might be acceptable. If typical workloads involve large scans, then more processing power is needed. Understand the types, frequency, plans, and typical concurrency for the application, ideally via prototyping the workloads and queries whenever possible.

Memory

When sizing an EsgynDB Enterprise Supermicro cluster for memory usage, keep in mind that many Hadoop ecosystem processes are Java processes. Due to memory efficiency optimizations for the JVM, there is a significant restriction just below 32GB. Crossing this threshold results in less usable memory because the internal representation of pointers changes in a way that consumes significantly more space. Plan for these processes to use a heap size of 16-32GB each for optimal performance on a large cluster. Reducing the memory for these components affects performance significantly, so do careful tuning and analysis before choosing a smaller value. The primary users of memory in the EsgynDB database engine servers based on query concurrency and workloads.

Storage

When sizing an EsgynDB Enterprise cluster for disk usage, for data nodes, SSD is only beneficial for high concurrency write. However, HDD drives are sufficient, HDD data disks configure disks as direct attached storage in a JBOD (Just a Bunch of Disks) configuration. RAID striping slows down HDFS and reduces concurrency and recoverability. For control nodes, data disks can be configured as either JBOD or RAID1 or RAID10.

As with processing power, disks are a unit of parallelism. For a given total-disk-per-node value, if workloads include many large scans, it is often most effective to have more smaller disks than fewer larger disks per node on data nodes. The reference architecture assumes that most workloads include large scans. HBase SNAPPY or GZ compression is strongly suggested. SNAPPY has less CPU overhead, but GZ compresses better. Degree of compression varies widely depending on the data and workload patterns, but generally accepted calculations suggest around a 30%-40% reduction, depending on data.

Compression adds to the path length for reading and writing, which can have an effect on data growth and ingest. Compression happens at the HBase file block level, limiting the amount of un-compression required at read time. When calculating overall disk space and data disk space per node, be sure to account for working space and anticipated ingest/outflow per node. Also remember that blocks of an HDFS file come with a replication factor (typically set to 3, so 3 copies of the data). That means that each 10 GB file occupies 30GB on disk. Esgyn recommends leaving approximately 33% of disk space free for overhead workspace.

Network

In general, 10GigE is the standard for networking for data traffic within an EsgynDB cluster. Using a slower network for data flow can significantly impact performance. 2 bonded 10GigE networks provide more throughput for I/O intensive applications.

Supermicro Server platform

The Supermicro TwinPro^2 shown in Figure 2 below, is an excellent choice as the server platform for the edge node, management node and head nodes.

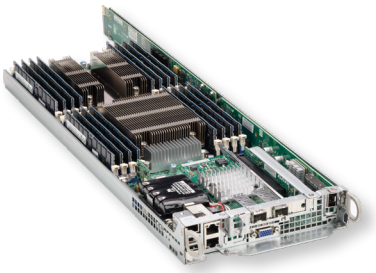

Figure 3. *1 of 4 dual-processor nodes (all pictures are for illustration purposes only)*

Figure 2. *Rear View of Supermicro 2U TwinPro2 SYS-2028TP-HC0R-SIOM*

Table 1. *Technical Specifications of Supermicro 2U TwinPro2 SYS-2028TP-HC0R-SIOM*

Processor Support	Intel® Xeon® processor E5-2600 v4/v3 family (up to 145W TDP)
Chipset	Intel® C612 chipset
Memory Support	Up to 2TB ECC 3DS LRDIMM , up to DDR4- 2400MHz ; 16 DIMM slots per node
Expansion	1 PCI-E 3.0 x16 (LP), 1 PCI-E 3.0 x8 (LP), 1 SIOM card support
Onboard Storage Controller	Broadcom 3008 SAS3 controller (8 ports); RAID 0, 1, 10
Drive Bays	6 Hot-swap 2.5" SAS/SATA drive bays
Connectivity	SIOM Networking
Power Supply	2000W Redundant Power Supplies Titanium Level (96%)

Network Configuration

Dual port 10G Ethernet network card: dual-port 10Gb Ethernet with SFP+ per node

Reference architectures

The following sections illustrate a reference progression of Hadoop clusters from a single-rack to a multi-rack configuration.

Single 2U TwinPro² Architecture

The single 2U Supermicro EsgynDB Enterprise Reference Architecture is designed to perform well as a single-rack cluster design but also form the basis for a much larger multi-2U design. When moving from the single-2U to Multi-2U design, one can simply add racks to the cluster without having to change any components within the single-2U.

Multiple 2U TwinPro² Architecture

The multi-2U design assumes the single unit design is already in place and extends its scalability. The single unit configuration ensures the required amount of management services are in place for large scale out. For multi-2U clusters, one simply adds extension racks of a similar configuration to the single-rack configuration.

Performance Notes

Supermicro EsgynDB appliance delivers performance, scale and stability, which enables customers to deploy Supermicro clusters with multi-2Us to run millions of queries, and at high concurrency. Standard benchmarks like YCSB, TPC-C and TPC-DS results showed that EsgynDB is the only mixed-workload engine that can consistently run OLTP, ODS, BI and analytical workloads.

Esgyn Cluster Tuning/Optimization

Optimizer technology represents one of EsgynDB's greatest sources of differentiation versus alternative SQL on-Hadoop projects or products. Two primary areas are the extensible nature of the optimizer to adapt to change and add improvements; and the sophistication and maturity level of the optimizer to choose the best optimized plan for execution.

About Super Micro Computer, Inc.

Supermicro® (NASDAQ: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide. Supermicro is committed to protecting the environment through its "We Keep IT Green®" initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

Learn more on www.supermicro.com

About Esgyn

Esgyn is the leader in Converged Big Data solutions that empower global enterprises to realize the potential of Big Data. With the industry's most mature, scalable and adaptive SQL Database for Big Data, Esgyn is leading the way enterprises cope with ever increasing data management needs on-premises and in the cloud.

United States of America (USA)

691 S. Milpitas Blvd., Suite 117
Milpitas, CA 95035 • +1-408-800-6916.
info@esgyn.com

Learn more on www.esgyn.com

Peoples Republic of China (PRC)

Beijing, Shanghai, Guiyang
PRC • +1-86-21-5082-2117

Learn more on www.esgyn.cn

Printed in USA ♻ Please Recycle ESGYN_180118_Rev09