

TRAVERSE: DISTRIBUTED, SCALABLE, HIGH-AVAILABILITY ARCHITECTURE

SUMMARY

Zyrion's Traverse is a breakthrough business service monitoring and network performance management application that provides real-time visibility into the health of IT services. Traverse provides a correlated and business-oriented view of the IT infrastructure, and links the underlying IT infrastructure to the supported business services and processes. Traverse's innovative Business Container technology enables IT and business personnel to create unique virtual views of the IT infrastructure. Traverse can monitor, in real-time, millions of metrics from tens of thousands of end points and devices.

Related white papers are available at:

<http://www.zyrion.com/company/whitepapers/>

- Virtualization and Private-Cloud Monitoring
- Business Service Management Technology for Connecting IT Infrastructure to Business Services
- Innovative Approaches for Identifying Root Causes of Service Performance Degradation

View online demo of advanced BSM software:

<http://www.zyrion.com/demo/>

Download free trial Zyrion Traverse:

<http://www.zyrion.com/download/>

THE DATABASE BOTTLENECK CHALLENGE

A key limitation of most management systems is the existence of a centralized database for processing of performance data. Even if the collection of data is managed by distributed components, the solutions invariably require centralization of the data for processing and alert generation. For large infrastructures, this introduces a significant performance bottleneck. The multiplier effect of the amount of data that needs to be processed as new devices are added is enormous. For each new device, hundreds of new metrics may need to be captured on a recurring frequency. For mission critical systems, the refresh frequency for performance data could be quite high. An environment of ten thousand devices, for example, may easily generate millions of metrics every few minutes. Capturing and processing these metrics in a single centralized database puts immense pressure on the overall application, creating a significant performance bottleneck.

FULLY DISTRIBUTED, REAL-TIME ARCHITECTURE

Traverse is built on a powerful, fully-distributed architecture featuring two primary components, Data Gathering Engines (DGE) and the Business Visibility Engine (BVE). What is unique about Traverse is that

there is no centralized data warehouse, unlike any other monitoring solutions that have to centralize their data to perform analysis, generate alarms and produce reports. Traverse has distributed collection capability AND a distributed database architecture, which allows the system to scale to extremely large environments with standard hardware.

Data Gathering Engines (DGEs):

Traverse's Data Gathering Engines are independent data gathering element that collect, process, analyze, and store performance data locally from networks, devices, servers and applications. DGE's collect the alarms, events, and polls from the various infrastructure elements such as servers, routers, and applications. Each DGE has its own distributed relational database. Unlike other architectures, there is no central database into which all DGEs have to feed data into for reporting purpose. Each DGE collects data according to the definitions it has been given from the object store. These definitions include which devices it should monitor, how frequently it should monitor such devices, and what it should do with the results.

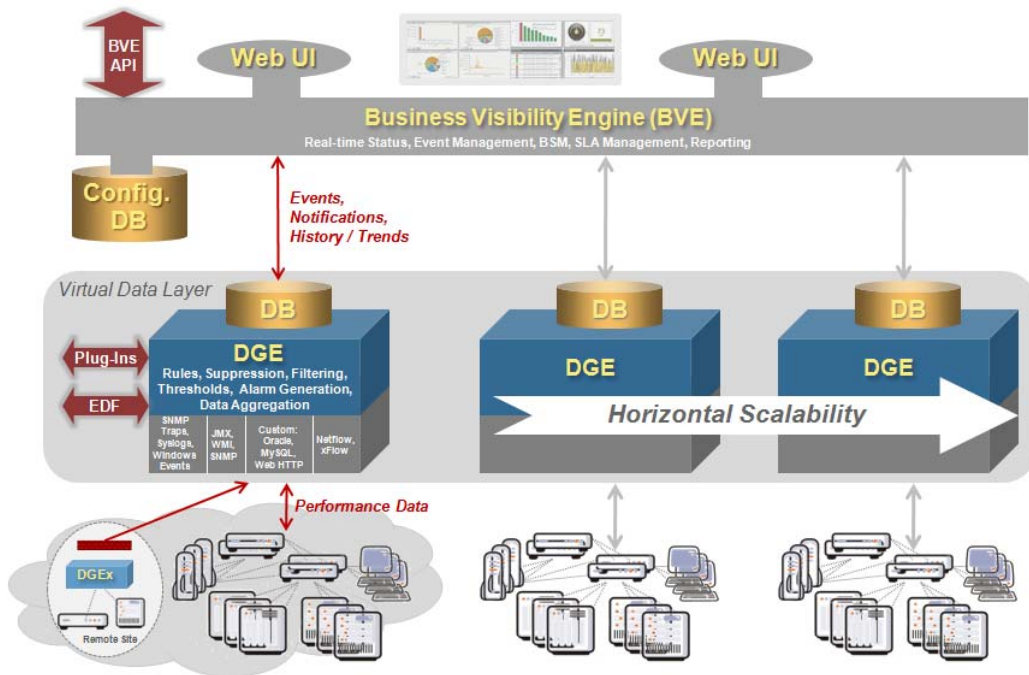


Fig 1: Enterprise and MSP Network Configuration

Traverse's Data-gathering Engine (DGE) Extensions enable capturing performance data from closed or secure networks by eliminating the need for inbound connections from the primary data aggregation point. Furthermore, the DGE Extensions do not require any static IP addresses, and thus new data gathering elements can be added quickly whenever and wherever needed. The DGE components also support monitoring multiple private networks that may have over-lapping IP addresses.

Business Visibility Engine (BVE):

The BVE is an intelligent configuration object store/repository that keeps track of users, permissions, devices, connections, and the correlation between services and underlying infrastructure. The object store allows IT personnel to set up business services and containers, and maintains the data on all thresholds, traps and syslogs. It also maintains application user permissions, definitions of what infrastructure elements are in which business container views, which data gathering elements monitor which infrastructure components, and at what frequency and with which specific queries/tests.

When a user logs in via the Web UI, and the data is automatically fetched from the distributed DGEs and presented in a unified, correlated view. Very little data is transmitted over the WAN. Traverse’s web-based reporting and graph generation engine consolidates system and application performance to provide real-time status and historical reports. Traverse BVE is responsible for coordinating information from all the DGEs and providing it to the user on-demand.

DISTRIBUTED DATABASES & PROCESSING

In a distributed environment, requests for reports and information can originate from a user via the web browser interface. A user may request, for example, a view of the 'Top 10' problems at that point in time. The request is received by the BVE software module which then fans out the request to the appropriate DGEs that are monitoring the elements that are of interest to the user who requested the report. Each DGE queried by the BVE will then run its Top 10 report and return the results to the BVE.

The BVE then summarizes the multiple results into a consolidated view and presents it to the user. This unique interaction between the BVE and the multiple DGEs is highly efficient. Zyrion has been awarded multiple patents on its distributed database architecture, and distributed processing and reporting solution. Most systems today require a central database for historical data reporting. This severely limits scalability. Zyrion's BVE removes the need for a central database through its coordination of multiple independent DGEs.

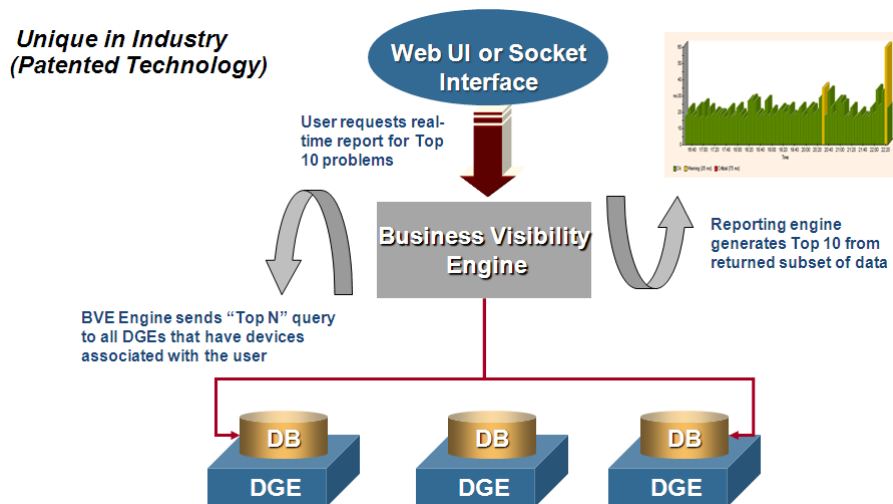


Fig 3: How Distributed Processing Works

HIGH AVAILABILITY CONFIGURATIONS

Traverse’s distributed database and processing architecture supports high levels of fault tolerance and scalability. All of the components in the various tiers are horizontally scalable which is essential for expansion and generation of real-time performance reports. All of the configuration information is stored in the BVE configuration database. On startup, the DGEs connect to the BVE configuration database and download a local copy of their configuration. Any updates made to the BVE configuration database are pushed out in real time to the corresponding DGE. If for some reason the BVE goes offline, the DGE can continue capturing and processing data, and generating alarms.

Warm Stand-by DGE:

To handle the case of a DGE physical server going down, a spare 'warm standby' server can be set up in any central location (N+1 redundancy) which has the software installed and configured. In the case of a production DGE going down for an extended period of time due to hardware failure, the name of the DGE can be set in the dge.xml configuration file and Traverse can be started up on the backup server. This backup DGE automatically connects to the BVE configuration database and downloads the configuration of the failed DGE. When the production DGE comes back up, it can run in parallel before shutting down the backup DGE. The only caveat is that the performance data collected during this interval will be missing on the production DGE.

If desired, a backup DGE can be installed for each of the production DGEs (N+N redundancy), but this is not really needed if the centralized DGE can poll all the data remotely. If connectivity between the DGE and the BVE database is lost, the DGE continues to poll, aggregate and even generate alarms completely independently. When connectivity to the BVE database is restored, the DGE restarts and synchronizes its configuration database.

Hot Replication of BVE:

The BVE database can be replicated on multiple servers for fault tolerance. DGEs are able to automatically fail-over to an alternate BVE database if the primary database is not reachable. The performance database which is local to each DGE can be located on a remote database cluster if needed for fault tolerance also. The JDBC communication between the DGE and the performance database allows creating such a setup seamlessly by making just a few configuration file changes.

Lastly, the Web application and reporting engine also gets all the configuration information from the BVE database server on startup and hence any number of Web application servers can exist behind a load balancer for fault tolerance as well as distributed report processing.

PROVEN SCALABILITY AND RELIABILITY

Hundreds of organizations are leveraging Zyrion’s technology for true IT business service assurance. Zyrion's customers include large and medium enterprises, MSPs, utilities, universities and governmental

institutions. Zyrion is delivering substantial benefits to customers in the areas of massive scalability, cost savings, flexibility, rich reporting and virtualization monitoring.



The table below illustrates the types of environments that are successfully being monitored and supported by Traverse.

Customer Name	Description	Environment Overview
Ancestry.com	Pioneer in online family history research and services, hosting 20 Million family trees containing 2 billion plus profiles	Traverse is monitoring over 10,000 servers and 2,000 networking devices (routers & switches) across 3 datacenters. Approximately 500,000 metrics are being processed every 5 minutes by the solution.
Sony Computer Entertainment of America	Leader in variety of entertainment products and service, with millions of customers and users worldwide	Traverse is monitoring a network of over 6,000 devices and applications using 20 Data Gathering Engines (DGEs) spanning 12 datacenters across Asia, Europe and North America. The virtual data layer within the solution is monitoring close to 1 million metrics every 5 minutes. The implementation utilizes 20+ different types of plug-in script tests, with close to 3,000 test instances.

THE BOTTOM LINE

Traverse has a distributed and scalable architecture that has the unique characteristic of not having a centralized data warehouse. Traverse has a distributed collection AND distributed database architecture which allows the system to scale to very large environments with standard hardware. The Data Gathering Engine’s (DGE) are data collectors – each collector can monitor between 500 to 1000 devices. The BVE layer is responsible for correlation and reporting, and as users log in via the Web UI, and the data is

automatically fetched from the distributed DGEs and presented in unified, correlated views. Traverse's architecture is inherently fault tolerant, and is well suited to monitoring mission critical, complex, distributed infrastructures.

ABOUT ZYRION, INC.

Zyrion is a provider of Business Service Management (BSM) and IT infrastructure monitoring software for enterprises, MSPs and governmental organizations. The company's business service container technology allows organizations to more easily and effectively manage IT-enabled business processes and services. Zyrion's flagship Traverse solution provides correlated, end-to-end network and server monitoring capabilities that link underlying applications and the IT infrastructure to business services. Zyrion has corporate offices in Sunnyvale, Calif. For more information, go to www.zyrion.com or call +1-877-7-ZYRION.