

Measurements made by eG Agents

Type	Details	Measurements
Operating Systems	Solaris 7.0 or higher, Red Hat Linux 6.0, AIX 4.3.3 or 5.x or 6.1 or 7, HP-UX 11, Free BSD 5.4, Tru64 5.1, Windows Vista, Windows 2008, Windows 7, Windows 8, Windows 10, Windows 2012, Windows 2016	<ul style="list-style-type: none"> System CPU and memory statistics such as CPU utilization, run queue length, blocked processes, swap memory in use, and available free memory; Disk statistics including percentage utilization of partitions, read and write rates on each physical disk, percentage busy, request queue length; Network traffic statistics including packet rate to and from the different network interfaces, bandwidth in and out of each of the interfaces (Windows only); TCP statistics such as incoming and outgoing connection rate, current connections, connection drop rate, etc.
	Operating systems that support the HOST Resources MIB	<ul style="list-style-type: none"> Monitoring of devices accessible via the server, which include the current status of a device and the number of errors encountered by the device; Monitoring of storage areas on the server, including tracking the storage area size, used space, and percentage utilization, etc.; System monitoring to determine the users currently logged in to the server and the total number of processes that are running; Tracking CPU usage on the server; Process monitoring including tracking whether a specific process is running or not and its memory utilization; TCP connection monitoring including number of currently established connections, failed connections, retransmission rate for segments, etc.
	AS400	<ul style="list-style-type: none"> Current workload indicators such as the number of jobs executing on the server; Metrics revealing the message processing ability of the server, which include, the number of messages in queue; Statistics that indicate the health of the storage subsystem, such as, the capacity of the auxiliary storage space and the percentage of space used, the amount of storage space used by temporary objects and permanent objects, etc.; User activity related metrics, such as, the number of users who have currently signed on, the number of user jobs that were temporarily suspended, the number of disconnected jobs, etc.

	Novell Netware	<ul style="list-style-type: none"> Monitoring of system memory available and usage, and utilization of memory by individual NLMs; Tracking of the Netware file system accesses (read and write rates), as well as monitoring all the volumes on the server to determine percentage utilization of the volume and the amount of freeable space on the volume; Tracking the Netware system processor usage and interrupt rates
	Apache, iPlanet/SunONE, Microsoft IIS, and Oracle HTTP web servers	<ul style="list-style-type: none"> Server availability and response time monitoring from an external location HTTP/HTTPS protocol emulation; Web server statistics including internal measures of connection and request rates handled by the server, percentage of errored responses, percentage aborts, etc. for the entire server; Reporting of similar metrics for web sites hosted on the server as well; Monitoring server processing times for individual transactions using eG's web adapter capability- – offering metrics such as request and connection rate for individual web transactions, response time monitoring per transaction, percentage aborts for individual transactions; Emulation of multi-step web service interactions using the eG client emulator to offer availability and response time statistics for each step of the service interaction
Web Servers	Sun Java Web Server 7	<ul style="list-style-type: none"> Metrics revealing bottlenecks in connection requests, such as, the number of connections in queue, the number of connections dropped from the queue, the duration for which connections were enqueued, etc. Metrics indicating whether/not the DNS cache has been enabled, and if so, how is it been utilized; Metrics measuring the file cache usage, such as, the number of cache hits/misses, the number of cache content hits/misses, the number of entries in cache, the cache heap size, whether the file cache is enabled/disabled, etc. Metrics revealing quality of the HTTP service on each virtual server, such as, the rate of HTTP requests to the virtual server, the number of errors that were logged, the data received and transmitted by the virtual server, the number and type of responses, etc. Metrics revealing the health of the web server instance, such as, the duration for which the instance was running, the rate of requests to the web server instance, the number of errors logged by the instance, the average response time of all virtual servers in instance during the last 5 minutes, the data received/transmitted by the instance, etc. Metrics indicating the uptime and session activity on each web application on the web server, which include, the number of JSPs loaded/reloaded on to each application, the number of active sessions, number of rejected sessions, the average lifetime of expired sessions, etc.

Web Servers		<ul style="list-style-type: none"> Metrics revealing how the JVM manages memory resources, such as, the heap memory size, the time taken for garbage collection, the number of times garbage collection occurred, the number of classes loaded/unloaded from JVM, etc. Metrics that indicate issues with keep-alive connections to the web server instance, such as, the number of successful requests from a keep-alive connection, the number of keep-alive connections that were flushed and refused, the number of keep-alive connection timeouts, etc. Metrics measuring thread pool usage, such as, the number of threads in pool, the number of requests queued, the number of idle threads in pool, etc. Metrics revealing bottlenecks in connection requests, such as, the number of connections in queue, the number of connections dropped from the queue, the duration for which connections were enqueued, etc.
	Nginx Servers	<ul style="list-style-type: none"> Metrics revealing the total number of connections, active connections, connections handled by the server etc, Measures indicating the total number of requests handled, the number of request header reads, total number of write responses, total number of waiting connections, average number of requests handled etc;
Middleware Application Servers	iPlanet Application Server	<ul style="list-style-type: none"> Server usage in terms of request rate, percentage of threads utilized, data transfer and reception rate; Server performance as measured by average response time
	WebLogic Application server	<ul style="list-style-type: none"> Server usage and performance measurements including throughput, percentage of heap space in use, number of requests queued; JDBC connection pool usage in terms of percentage of connections used, and percentage of connections pending per pool; Statistics pertaining to the EJB components such as transaction commit and rollback rates, timeouts during access to individual components, access queue lengths, etc.
	ColdFusion Application server	<ul style="list-style-type: none"> Database measures including database access rate and average response time for database access; Server workload measures such as request rate, data transmission and reception rate, queue length, request time out rate, and average response time
	SilverStream Application server	<ul style="list-style-type: none"> Statistics such as the utilization of the server, percentage of threads associated with the client connections, the percentage of idle threads, memory usage, etc.; Server workload measures such as request rate, response time, transmission rate, the number of sessions, etc.

Middleware Application Servers	WebSphere Application Server	<ul style="list-style-type: none"> • Internal measures of server availability and memory usage; • Metrics indicating cache usage such as the number of cache hits, misses, etc. • Measures that help determine the sizing needs of the JVM heap, such as the percentage of JVM memory utilized, the free memory, the number of GC executions on the heap, etc. • Comprehensive monitoring of thread and connection pools including measurement of pool utilization, wait counts per pool, average wait time for each pool, and creation and destruction rates per pool; • Measurements of individual EJB components such as the rate of instantiations, average creation time and response time per EJB component, etc.; • Metrics associated with global and local transactions such as the number of transactions begun at the server, number of concurrently active transactions, average duration of the transactions, etc.
	JBoss Server	<ul style="list-style-type: none"> • Measures JVM heap performance related to used & accessible memory in JVM Heap. • Monitors thread pool usage, connection pool usage, and alerts when pool usage levels reach threshold limits. • Monitors Servlet executions, so the most popular servlets can be chosen for further optimization; Also tracks EJB pools and the activation/destruction of the beans.
	JBoss AS/EAP servers	<ul style="list-style-type: none"> • Measures such as the incoming requests processed, time taken for processing the requests, errors encountered etc, relating to each JBoss connector; • Measures such as the active connections, available connections, connections destroyed, connections created, time taken for connection creation etc can be obtained by extensively monitoring the JBoss datasource apart from obtaining metrics relating to the connection pool; • Measures relating to the prepared statements of the datasource such as the prepared statement cache access size, number of times the cache was accessed etc., • Critical metrics relating to the EJBs such as the number of times the EJB has been invoked, bean instances that are available in the pool etc., • Measures revealing the size of the MQ Queues and the type of messages processed by the queue etc., • Measures revealing the messages in the MQ topic, the message processing rate, subscribers count for the topic etc., • Metrics revealing the requests processing capability of each servlet deployed in the JBoss AS 7;

		<ul style="list-style-type: none"> • The transaction processing capability of the Jboss AS 7 and the numerical statistics pertaining to each type of transaction; • Measures such as the active connections, available connections, connections destroyed, connections created, time taken for connection creation etc., can be obtained by extensively monitoring the JBoss XA-datasource apart from obtaining metrics relating to the connection pool; • Measures relating to the prepared statements of the XA-datasource such as the prepared statement cache access size, number of times the cache was accessed etc.,
Middleware Application Servers	JRun Application Server	<ul style="list-style-type: none"> • Measures pertaining to the thread pool usage of all the server instances of a JRun server such as the percentage of threads utilized, threads queued etc.; • Thread processing metrics such as the waiting time of a request for processing, processing time of the server and a request, the data reception and transfer rate of the server, etc.; • Performance statistics related to the JVM of all instances of a JRun server, such as the memory used by the server, free memory, number of active sessions in the server and the server memory
	Borland Enterprise Server (BES)	<ul style="list-style-type: none"> • Key runtime statistics pertaining to a BES management agent, such as the amount of memory available to and consumed by the agent, the current number of active threads in the agent, etc.; • Partition-specific statistics such as the memory available to and consumed by the partition, the current number of active threads in the partition, etc.; • Performance statistics related to the BES EJB container, which includes the time spent in performing CMP related activities such as executing SQL queries and updates, the time spent in receiving and serving a TCP request, in dispatching methods to various objects, time spent in the ORB, etc.; • Important statistics measuring the time taken for performing JDBC1 and JDBC2 related activities, such as the time spent by the data sources in acquiring a pool connection, in registering a transaction resource in transaction service, etc.; • Measurements related to the time spent in activating and passivating a stateful session bean; • Transaction-specific statistics such as the time spent in beginning, committing, and rolling back transactions

Middleware Application Servers	SunONE Application Server	<ul style="list-style-type: none"> • Overall health of the virtual server in terms of rate of requests to the web server, data reception and transfer rate, etc.; • Metrics related to JDBC connection pools such as the number of threads awaiting the database connection, the number of failed connection attempts, etc.; • Transaction-related measures including the number of transactions that are completed, rolled back, and in-progress; • EJB caching and pool-related metrics, which include the number of beans in a pool, the number of times user requests found and failed to find a bean, etc.
	Orion Server	<ul style="list-style-type: none"> • Measures pertaining to the JVM running on an Orion server which includes the active thread count, the total memory capacity of the JVM, and the unused memory in the JVM
	Tomcat Server	<ul style="list-style-type: none"> • Reports JVM details specific to loading /unloading of classes, checks thread pool usage, alerts on deadlocks or excessive non-daemon threads that could be impacting the server's performance. • Monitoring the GC (Garbage Collection) process, reports on memory usage patterns in the JVM. • Measures the workload to the Tomcat connector. • Measures JSP engine performance and reports on usage and execution of JSPs and servlets.
	Oracle 9i/10G Application Server	<ul style="list-style-type: none"> • Overall JVM performance in terms of JVM availability, current active threads in the JVM, the high and low watermarks of heap usage, etc.; • JDBC-connection related metrics such as the number of threads creating connections, time taken for connection creation, etc.; • Transaction-related metrics such as the number of active transactions, transaction commits, and rollbacks; • Web module-specific metrics such as the number of threads handling requests, the number of completed requests, request processing time, etc.; • JSP-related metrics including the number of active requests for the JSP, request processing time, etc.; • Servlet-related metrics such as the number of threads servicing the servlet, the number of completed requests, request processing time, etc.
	Oracle Forms Server (Windows)	<ul style="list-style-type: none"> • Key metrics reporting the memory and CPU usage of the server; • Metrics related to user sessions such as the number of active and idle sessions, session duration etc.; • Measures that report server responsiveness such as the rate of requests received, response time of the Forms server and the database server, etc.; • General statistics pertaining to database users that includes such as the rate of requests received from a user, response time of the Forms server and the database server, memory and CPU utilized by the user

Middleware Application Servers	Glassfish Monitoring <ul style="list-style-type: none"> Key measures of the processing ability of the applications deployed on the server, such as, the number of requests to the server per application, the processing time of requests per application, etc.; Session-related metrics, such as, the number of active sessions, the number of activated and passivated sessions, expired sessions, rejected sessions, etc., Thread pool usage metrics such as, the number of core threads in each pool, the number of busy threads, the maximum number of threads allowed per pool, the percentage usage of threads per pool, etc. Metrics pertaining to transactions, such as the transaction state and the number of active/committed/rolled back transactions; Connection pool usage metrics, which include, the number of free and used connections per pool, the waiting time of connections, potential connection leaks, the length of the queue of waiting connections, failed connections, etc.; Statistics indicating the processing ability of servlets/servlet groups, which include, the number of requests processed by a servlet/servlet group, the average processing time of requests, the count of servlet errors, etc.; Metrics revealing how well the EJB cache has been utilized, which include, the number of cache hits/misses, the cache hit ratio, etc.; Metrics providing method-level performance insights, such as, the number of successful/error-ridden invocations of a method, method execution time, etc.; EJB pool-related metrics, such as, the count of beans in a pool, the number of threads waiting for a free bean in a pool, the count of beans created/destroyed in a pool, etc.; Measures revealing the level of activity on stateful and stateless session bean containers, such as, the number of times create, ready, and remove method calls were invoked on these containers.
	Domino Application Server <ul style="list-style-type: none"> Statistics measuring the health of the web server component of the Domino application server, such as, the load on the server in terms of requests to the server, the data received by it, and the count of current connections to it; Replicator-related metrics such as, the number of successful and failed replications, the number of pending replications, etc. Statistics pertaining to the Domino database such as number of databases in cache, the percentage of pages cached, the maximum capacity of the server cache; Memory usage metrics, which include the total memory allocation, the shared memory allocation, etc.; Important metrics pertaining to the network traffic, that includes, the data receipt and transfer rates, etc.

Middleware Application Servers	Tuxedo Domain Server	<ul style="list-style-type: none"> • Measures indicating the availability and response time of a Tuxedo application; • Domain usage statistics such as the percentage of servers in the domain that have been utilized, etc.; • Queue-related measures including the queue state, servers associated with the queue, etc.; • Site and server-related measures, such as the site state, server state, the number of current logins, bridge state, transaction commit rate, etc. • Connections that are currently open on the server; • Bulletin Board related statistics such as the servers registered on the bulletin board, services, interfaces and queues available in the domain, etc,
	WildFly JBoss	<ul style="list-style-type: none"> • Measures such as the incoming requests processed, time taken for processing the requests, errors encountered etc, relating to each connector; • Measures such as the active connections, available connections, connections destroyed, connections created, time taken for connection creation etc can be obtained by extensively monitoring the datasource apart from obtaining metrics relating to the connection pool; • Measures relating to the prepared statements of the datasource such as the prepared statement cache access size, number of times the cache was accessed etc., • Critical metrics relating to the EJBs such as the number of times the EJB has been invoked, bean instances that are available in the pool etc., • Measures revealing the size of the MQ Queues and the type of messages processed by the queue etc., • Measures revealing the messages in the MQ topic, the message processing rate, subscribers count for the topic etc., • Metrics revealing the requests processing capability of each servlet deployed in the WildFly; • The transaction processing capability of the WildFly and the numerical statistics pertaining to each type of transaction; • Measures such as the active connections, available connections, connections destroyed, connections created, time taken for connection creation etc., can be obtained by extensively monitoring the XA-datasource apart from obtaining metrics relating to the connection pool; • Measures relating to the prepared statements of the XA-datasource such as the prepared statement cache access size, number of times the cache was accessed etc.,
Database Servers	Oracle database	<ul style="list-style-type: none"> • Overall health of the server in terms of availability and response time; • Key metrics of database health including library cache hit ratio, data buffer cache hit ratio, dictionary cache hit ratio, redo log buffer miss ratio, percentage of sorts on the disk, monitoring of hot datafiles, tablespace usage, user sessions, lock and latch activity, top SQL queries in terms of resource usage, etc.

Database Servers	Oracle RAC	<ul style="list-style-type: none"> External metrics indicating the availability and responsiveness of the cluster; Measures revealing the session load on the cluster; Key metrics pertaining to the number of sessions that are waiting for a redo log write confirmation after a commit; MTTR-related statistics, which include, the target and estimated MTTR, target and actual number of redo blocks, etc.; Metrics related to wait events, such as, the number of events of type log file parallel write, Db file parallel write, log file sync, and Db file sequential read that occurred on the cluster, the number and nature of wait events that occurred per session module, the number and nature of session wait events, etc.;
	Rocket UniVerse Database Server	<ul style="list-style-type: none"> Measures revealing the total number of sessions, percentage of sessions used by the user, and number of interactive and background processes initiated by the user. Metrics related to number of active group locks and active file/record locks. Metrics revealing undo tablespace usage such as, the number of sessions accessing the undo tablespace and their usage duration; top sessions to the undo tablespace in terms of duration of usage are also revealed; Locking activity-related metrics, which include, number and duration of transaction locks; Tablespace usage metrics, which include, the size of tables per tablespace, index size, partition size, used space, free space, etc. Key metrics related to the queries executing on the undo tablespace, such as, the query duration, time taken for undo retention, etc.;
	DB2 UDB	<ul style="list-style-type: none"> Instance level measures such as the percentage of agents waiting, percentage of piped sorts, the number of registered agents; Statistics revealing how effectively the agent pool was utilized, which includes, the percentage of agents that were directly assigned from the pool for servicing requests, the number of agents stolen from an application, etc.; Key metrics of database health including lock activity, buffer pool hit ratio, catalog cache hit ratio, read and write rates, rollback rate, transaction rate, average sort time, percentage of sorts overflowing, heap space used by each sort, SQL activity, etc.; Lock activity related measures such as the number of locks currently held by all applications on the database, the rate of lock waits, the duration of lock waits, the number of deadlocks, etc.; Buffer pool usage metrics, such as, the percentage of pages, data pages, and index pages readily served by the pool, the percentage of times dirty pages and victim pages were cleaned from the database, etc. External measures such as the availability and response time of the database server

Database Servers	DB2 DPF	<ul style="list-style-type: none"> • Metrics revealing how effectively the agent pool has been utilized, such as, the total number of agents in pool, the number of idle agents, percentage of pool agents used for servicing requests, etc. • Metrics revealing the database manager's health, which include, the rate at which sorts requested heaps post threshold, the rate of piped sort requests and rejects, etc. • Locking related metrics, such as, the number of locks held, escalated, timedout, the number and duration of locks waits, etc. • Buffer pool usage metrics, which include, the hit ratio, prefetc.h ratio, catalog cache hit ratio, package cache hit ratio, etc. • Sorting related metrics, which include, the sort heap space allocated to each partition, sort heap space used, average sort time, etc. • Statistics measuring the level of I/O activity on the partitions, such as, direct read rate, direct write rate, etc. • Key metrics pertaining to the transactions executing on each partition, such as, rate of transactions, commit rate, and rollback rate
	Sybase Adaptive Database	<ul style="list-style-type: none"> • Overall health of the server in terms of availability and response time; • Key metrics of database health including lock activity, the number of network bytes received and sent, cache utilization and cache hit ratio (of both the data cache and procedure cache), the rate of reads and writes performed on a device, space usage of the SQL databases, transaction related metrics, etc.
	Microsoft SQL	<ul style="list-style-type: none"> • External monitoring of server availability and response time for a typical query; • Key metrics of SQL server health including various caching hit ratios, page read and write rates, pending transactions, locking rates and wait time, memory usage of various components such as optimizer, locking subsystem, usage of the server's buffer memory, space usage of the SQL databases, etc.; • Monitoring of the server CPU usage, tracking number of running processes in the SQL server, and analysis of what queries the running processes are executing
	MySQL	<ul style="list-style-type: none"> • External measures of availability and responsiveness; • Rate of data transmission and reception by the server, number of active, aborted, failed connections, rate of queries handled by the server, and rate of slow queries, rate of commits and rollbacks, locking activity including number of locks and the percentage of lock waits

Database Servers	PostgreSQL	<ul style="list-style-type: none"> Statistics revealing table I/O, such as, the rate at which heap blocks and index blocks of a table are read, the percentage of heap blocks and index blocks of a table read from buffer cache, etc.,; Statistics revealing index I/O, such as, the rate at which index blocks are read, the percentage of index blocks read from buffer cache; The size of each tablespace on a PostgreSQL database; Metrics measuring background I/O, such as, the number of checkpoint requests received, the number of buffers released for re-use from the buffer cache, etc. Database usage related metrics, such as, the current database size, the percentage of requests to database serviced by cache, the rate of inserts, deletes, updates, rollbacks on database, etc.; Metrics measuring the rate of index scans, the rate at which rows are read/fetched from each index, etc.; Statistics on table usage, which include, the number of sequential scans and index scans initiated on a table, the number of times the table was scanned, the rate of inserts, updates, deletes on the table, etc.; User-related metrics, such as the number of connections established by a user on the server, the number of active, waiting, idle connections per user, etc.; Metrics revealing the availability of availability zones and the default region; Metrics reporting the number of locks currently held on the server; Metrics revealing the availability and responsiveness of the server; Metrics revealing the number and details of the long running queries on the server;
	Progress Database	<ul style="list-style-type: none"> Measures revealing the total number of sessions of each user and the rate at which the sessions are created by the user; Measures such as the availability of the database and the connection, the time taken by the database and the database connection to respond to the user queries, whether the query is executed successfully or not and the number of records that were fetched from the database; Metrics relating to the rate at which after-image writes are performed by the after-image writer, the total number of the buffer busy waits, etc; Metrics revealing the rate at which before-image writes are performed by the before-image writer, the before-image writes hit ratio, the total number of the buffer busy waits, etc; Statistics indicating the rate at which the read and write operations were performed on each data file, the rate at which the buffered read and write operations were performed on the data file, the rate at which the unbuffered read and write operations were performed on each data file, etc; Critical metrics revealing the number of index entries that were newly created and were deleted per second, the rate at which the index entries were scanned, the number of block splits that were created in the database while adding new index entries, etc;

Database Servers		<ul style="list-style-type: none"> Measures such as the rate at which each user accessed the database, the number of before-image and after-image read operations performed by each user per second, and the number of before-image and after-image write operations performed by each user per second; Metrics reporting the number of client requests received per second for performing the read and write operations on the database, the rate at which the read and write operations performed on the database and percentage of the data blocks that is readily available in the memory without doing disk I/O; Buffer related measures such as the total number of buffers in the cache, the number of buffers that are currently in use and are modified, the percentage of the used buffers and modified buffers, the number of buffers in the checkpoint queue, etc; Measures revealing the number of user sessions and the rate at which the number of sessions were created; Metrics revealing the rate at which the spin latch was created, the average time taken by the spin latch, etc; Measures relating to the lock such as the number of exclusive locks acquired by the user per second, the number of shred locks that were released per second, etc;
	Informix Dynamic Server	<ul style="list-style-type: none"> External measures of availability and response time; Session-related metrics such as the number of current and blocked sessions; transaction-related metrics such as the commit and rollback rates; efficiency metrics like rate of sequential scans, disk sorts, etc.; lock activity monitoring including tracking the number of lock requests, waits, timeouts, and deadlocks; space monitoring for each database
	Intersystems Cache	<ul style="list-style-type: none"> Reports buffer pool usage metrics such as the current buffer size, the number of interactive and batch buffers currently available, etc.; Measures the availability and responsive of the server; Reports critical database performance indicators such as lines executed; routine loads, new global references, new global sets, new global kills, logical database block reads, physical database global read and write, new database journal entries, write update status, etc.; Reports the current number of resource seizes. Monitors the users to the database instance, the database and routine caching activities performed by the database instance, license usage of the database instance, errors (if any) that have been logged, etc.; Monitors the caching and data management functions performed by the ECP application server; Reports statistics that reveal whether critical application and system processes are running or not;

		<ul style="list-style-type: none"> Measures the health of the locking activity by reporting the number of locks currently held; Monitors the Cache console log to reveal the number of warnings, and severe/fatal errors encountered by the Cache database
Database Servers	SAP HANA	<ul style="list-style-type: none"> Measures revealing the space usage statistics of the disk volume such as the total size, unused size, percentage of free space currently available, percentage growth in space usage; Measures revealing the percentage of CPU usage of the database server; the CPU time spent in system level processing, user level processing; percentage of time the CPU of the database server was idle etc., Memory related metrics such as the total physical memory, the free physical memory, physical memory that is currently used by the database server; total swap memory, used swap memory, swap memory that is currently used by the database server etc., Space related metrics for each service in the volume such as the used space, percentage of space that is currently used, rate at which the volume file has grown etc., Measures revealing the total size of each cache, numerical statistics revealing the requests or transactions available for retrieval from the cache, the inserts, invalidate requests etc., Measures indicating the number of expensive SQL statements and the time taken to execute the expensive statements etc., Measures indicating the uptime of the database, the workload on each service port of the database, the numerical statistics relating to the compilations performed, the records updated, the rolled back transactions through each port etc., Measures revealing the number of open transactions of each type and the maximum time taken for a transaction to execute; number of blocked transactions of each type; Measures revealing the availability and responsiveness of the database server; Metrics revealing the number of open connections, idle connections running connections etc for each user of the database server; Measures revealing the active status of each service of the database server; the uptime; average time taken by each service to respond to requests from clients; number of threads for each service etc.; Metrics related to the CPU usage of each service; memory usage statistics such as available memory, memory that is available for use etc.,
Access Applications	Microsoft Terminal Server	<ul style="list-style-type: none"> Monitoring of the server's TCP port for availability and response time; Tracking of the server's CPU and memory usage; Monitoring of specific applications deployed on the server to determine which applications are running, which are the most resource intensive applications, and who is using these applications; Session-related metrics to determine how many concurrent sessions are there – currently and at peak times; In-depth analysis of the applications invoked/used by individual users.

Access Applications	Citrix Web Interface	<ul style="list-style-type: none"> Key statistics revealing the health of the interactions between the web interface, XML service, and IMA service, which include availability of the Citrix XML service, the time taken to establish a TCP connection to the Citrix XML service, the total time taken for a user to login to the Citrix web interface and enumerate all the applications, etc.
	Citrix StoreFront	<ul style="list-style-type: none"> Metrics revealing the rate at which resources were accessed from the store, how well the ICA protocol and RADE process is utilized in accessing the store, time taken to access the resources using RADE process etc., Key metrics indicating the rate at which users are authenticated based on their language preference, average time taken for such authentication, how well the password change requests are processed for the users? And the time taken to change the passwords etc., Metrics revealing how well the authentication store stores the user requests, how well the user requests are deleted after the service has been serviced etc., Metrics related to Citrix Dazzle such as the rate of image response received for the resources accessed, rate at which the resources were accessed, arte at which cache calls were updated upon user requests etc., Measures revealing the user subscription details such as the rate at which user subscriptions were added, enabled removed etc from the store, the rate at which user subscriptions were updated etc., Metrics revealing the rate at which the users accessing through the Citrix Self service plugin are authenticated to access the controller and the average time taken for access;
	Citrix XenApp Server 4/5/6	<ul style="list-style-type: none"> ICA protocol level monitoring of Citrix server availability and response time; True client emulation by recording and playback of user transactions to measure each step of the user transaction – including login time, time for checking all published applications, time for accessing each published application, etc.; Application resource usage metrics, such as the number of instances of each application currently executing, and the CPU and memory utilization of the application, so as to allow administrators to determine the most resource intensive applications; User-related statistics such as the number of users currently logged in, their CPU and memory utilization, as well as reporting on what processes the user is currently executing; Session-related statistics that provide the number of active and inactive sessions, and the login times for each session; MetaFrame XP-specific performance metrics indicating the rate of application enumeration and resolution, the rate at which data has been read from and written into the IMA data store, availability of the data store, etc.; Monitoring at the Citrix farm level to determine the number of pool licenses that are in use, assigned pool licenses, etc.; Comparison of application load and license usage for each application across each of the servers in a farm; Monitoring of the license usage of the products of the Citrix MetaFrame Access Suite

Access Applications	Citrix XenApp Server	<ul style="list-style-type: none"> • TCP and ICA connection availability of each TCP port and its responsiveness; • Application resource usage metrics, such as the number of instances of each application currently executing, and the CPU and memory utilization of the application, so as to allow administrators to determine the most resource intensive applications; • User-related statistics such as the number of users currently logged in, their CPU and memory utilization, as well as reporting on what processes the user is currently executing; number of times the sessions were disconnected; users who reconnected soon after a disconnect, average client latency of each user., • Session-related statistics that provide the number of active and inactive sessions, connected and disconnected sessions, number of total sessions and the login times for each session;
	Citrix XenDesktop Director	<ul style="list-style-type: none"> • Key metrics revealing the numerical statistics of desktop OS machines in each delivery group configured in a site such as the total machines, machines in maintenance mode, Powered off machines, assigned machines etc; • Statistics revealing the number of machines of each type that are currently in the state of failure; • Metrics revealing the session load on each server OS machine, the resource usage of each machine and the current state of each machine; • Metrics revealing the user connections to each delivery group, average time taken for users to access the desktops/applications delivered by each group, number of logins that are slow etc; • Session-related statistics that provide the number of active and inactive sessions, and the login times for each session to each delivery group in a site, status of sessions to each delivery group, etc; • User connections related metrics such as the failure of connections due to client side problems, connection failure due to configuration errors, machine failures, number of new connections, login duration of each user, user authentication time, real time login details of each user etc;
	Citrix Access Gateway	<ul style="list-style-type: none"> • Key metrics related to the processing ability of the Citrix Access Gateway, which includes, the number of context requests pending processing, the rate at which the gateway performs commits, updates, deletes, etc., the rate at which data streams were read/written, etc. • Statistics pertaining to session management, such as, session cache hits and misses
	Citrix Netscaler LB	<ul style="list-style-type: none"> • Metrics revealing the extent to which CPU and memory resources were utilized by the Netscaler device • Network traffic-related statistics indicating the amount of data/packets transmitted and received over every VLAN configured on the device, and the packet loss ratio

Access Applications	Quality Virtual Desktop	<ul style="list-style-type: none"> Metrics revealing the current state of the server node and the blocked status of the server node, the number of virtual desktops in the node, successful/failed authentications, HTTP requests handled, metrics related to NX protocol attempts etc; Numerical statistics revealing the total hosts, powered on hosts, powered off hosts, starting hosts blocked hosts in the node cluster, disk images available in the server node etc; Desktop related statistics in the server node such as the registered desktops, powered on desktops, starting desktops, powered off desktops, desktops with users logged in, desktops without users etc; Key metrics relating to the currently active user sessions, new user logins, sessions that were logged out etc; Metrics revealing the total memory of the operating system flavor, overlay enabled status of the Operating system flavor, network availability, network delay of the server etc; A host of inside view metrics such as the percent of time the disk was busy processing requests, average time taken for reading the data from the disk, throughput of the virtual disk, total capacity of the disk, percentage of space utilized by the disk, physical memory utilization of the disk, incoming and outgoing traffic through each virtual desktop, CPU and memory usage of each virtual guest, TCP connections of the virtual desktop, TCP protocol traffic to and from each guest, latencies seen by users connecting to a virtual desktop, uptime of each virtual desktop etc;
	Oracle VM Server	<ul style="list-style-type: none"> Metrics revealing the number of processors, number of cores per socket, CPU utilization percentage; processors assigned to the control domain, the CPU utilization of the control domain etc; Memory utilization metrics such as the total physical memory, used physical memory, memory consumed by VMs etc; Desktop related statistics in the server node such as the registered desktops, powered on desktops, starting desktops, powered off desktops, desktops with users logged in, desktops without users etc; Key metrics relating to the capacity, space utilization, sharing status of the file system and the SAN storage, virtual CPU utilization of the control domain, swap memory and free memory available in the console, disk space utilization of the console, throughput of the disks available in the console etc; Metrics revealing the current status of the server, availability and responsiveness of the server;

		<ul style="list-style-type: none"> Statistics revealing the current operational state of each VM in the server, capacity of each disk, memory allocated and utilized by each disk, number of VMs registered, added, removed from the network etc; A host of inside view metrics such as the percent of time the disk was busy processing requests, average time taken for reading the data from the disk, throughput of the virtual disk, total capacity of the disk, percentage of space utilized by the disk, physical memory utilization of the disk, incoming and outgoing traffic through each VM, CPU and memory usage of each virtual guest, TCP connections of the VM, TCP protocol traffic to and from each guest, latencies seen by users connecting to a VM, uptime of each VM, disk partition alignment status etc;
Access Applications	Citrix Netscaler VPX/MPX	<ul style="list-style-type: none"> External metrics of availability and responsiveness of the device; Metrics revealing the overall health of core hardware components of the Netscaler device, such as, the current voltage and temperature of CPU cores, the speed of CPU fans and system fans, the temperature of CPU cores, the current voltage output of various power supply units, the current state of power supply units, etc.; Metrics revealing the CPU usage per processor and for the Netscaler device as a whole; Disk space usage metrics, such as, the total capacity per disk partition, the amount of used and free space, and the percent disk space usage per partition; Memory usage metrics revealing the percent memory used globally by the device and the per feature memory usage; Metrics revealing the current state of each network interface configured on the device, the link state of the interfaces, link uptime/downtime, data/packets received/transmitted over the interfaces, packets dropped by each interface, error packets received/transmitted by each interface, etc.; Metrics measuring load imposed by RNAT sessions, such as, the number of active sessions, the amount of data transmitted/received, etc.; Statistics monitoring VLAN traffic, which include, data/packets received/transmitted per VLAN, count of packets dropped by each VLAN, etc.; Measures indicating how well the device handles IP, HTTP, ICMP, TCP, and UDP load; Key metrics revealing the device's DNS query handling capability, such as, the number of queries received and answered, the number of times the cache was flushed, the count of queries for which records were not found, number of queries in invalid query format, the number of responses in invalid format, the number of DNS requests refused, etc.; SSL acceleration-related metrics, such as, number of SSL crypto cards present in the device, the number of SSL cards that are UP, the status of the SSL engine, the number of SSL/SSL v2/SSL v3 transactions performed on the device, the count of Front-End and Back-End SSL sessions, etc;

Access Applications		<ul style="list-style-type: none"> • Application firewall-related metrics, such as, amount of data received/transmitted by the firewall, the number of aborts, redirects, start URL and deny URL security check violations, etc.; • Measures pertaining to the High Availability configuration of Netscaler, which include, whether the Netscaler appliance is HA-enabled or not, the current state of the HA node, the number of heartbeat packets received and sent, the number of times command propagation timed out, the number of synchronization failures, etc.; • Load-balancing virtual server-related metrics, such as, the current state of each virtual server, the number of client and server connections to each server, the data/packet traffic on each server, etc.; • Metrics indicating the current state of and data/packet traffic handled by each authentication virtual server; • Statistics that reporting anomalies in transmission of log messages to the SYSLOG server, such as, the total number of log messages generated, number of logs sent and not sent to the syslog server, the count of NAT allocation failures, NSB allocation failures, NAT lookup failures, etc.; • Service-related metrics, such as, the current state of the virtual server bound to each service, the number of server and client connections to each service, the current load on the service, etc.; • VPN session-related metrics, such as, the number of VPN login requests received, the number of times the VPN login page failed, the number of DNS and WINS queries resolved by the VPN sessions, etc.
	Sun Ray Server	<ul style="list-style-type: none"> • Measures that indicate whether every Sun Ray client is able to connect to the Sun Ray server; • Session-related statistics revealing the number and type of sessions on the server; • Metrics that shed light on the current status of critical device services on the server;
	Citrix Netscaler HDX Insight	<ul style="list-style-type: none"> • License usage metrics such as the number of SSL VPN licenses configured, number/percentage of licenses currently utilized, etc. • Measures that indicate the number of user sessions that are currently active across all virtual desktops, number and percentage of new logins and number of sessions that were logged out. • Metrics revealing average rate at which data is transferred over the ICA sessions, number of unique users and number of sessions. • Key statistics revealing number of desktop sessions currently launched by each user, average latency experienced by each user due to problems with client and server side networks, screen lag experienced by this user, etc; • The average latency experienced by each session caused by the NetScaler appliance when ICA traffic flown from server network to client network and vice versa, host delay, etc;

Access Applications		<ul style="list-style-type: none"> Key metrics revealing the number of instances/sessions of each application that are currently launched, time taken to launch each application etc; Measures such as the rate at which data is transferred over the virtual channel, the total number of unique user sessions handled and unique applications launched by the Netscaler appliance etc; Metrics that reveal the number of user sessions that are currently active across all XenApp servers, number and percentage of new logins and number of sessions that were logged out. Metrics revealing the number of applications launched by each user, screen lag experienced by the user while interacting with an application on XenApp, screen lag time of the user smoothed over the client and server side connections, etc; Measures indicating how many times in each session the client and server advertised a zero TCP window, how many times the retransmit timeout got invoked in the session on the client and server side connections, etc.; The rate at which data is transferred over the ICA sessions launched from each device type that interacts with the applications/virtual desktops.
	Citrix Netscaler Web Insight	<ul style="list-style-type: none"> Measures revealing the number of requests and total amount of data received by each web server, etc; Metrics indicating the number of requests and total amount of data received by each web application, elapsed time between the end of an enquiry and the beginning of a response from the web application, average latency caused by the server and client networks, etc.; Measures revealing the number of requests and total amount of data received by the appliance; the number of requests and total amount of data received by the appliance using the HTTP request method etc; Measure indicating the elapsed time, from when the browser starts to receive the first byte of a response of specific type until either all page content has been rendered or the page load action has timed out. Measures revealing the number of requests received from clients running specific operating system, the number of requests and amount of data received from each type of client etc; Key statistics mentioning the render time and load time of each URL, number of requests received for this URL etc; Metrics revealing the number of requests received from each client, latency caused by the client-side network etc;
License Servers	Citrix License Server	<ul style="list-style-type: none"> License usage metrics such as the number of licenses installed on the server, the number and percentage of licenses utilized, etc.
	Terminal Services Licensing Server	<ul style="list-style-type: none"> License usage metrics such as the number of licenses installed on the server, the number and percentage of licenses utilized, etc.

Directory Servers	Netscape Directory	<ul style="list-style-type: none"> Key metrics revealing the performance of the caching activity of the server, which includes the cache hit ratio, rate of page reads and writes, etc.; Metrics on entry caches such as the cache hit ratio; External measures of availability, responsiveness, TCP connectivity, etc.
	Active Directory	<ul style="list-style-type: none"> Overall health of the Active Directory in terms of schema cache hit ratio, the number of pending requests, the rate of write operations on the directory store, etc.; Key metrics pertaining to the replication activity like the number of replication updates happening on the current domain controller, the number of changes applied to the object properties through inbound replications, etc.
	SunONE Directory	<ul style="list-style-type: none"> Database cache usage-related statistics, such as, the percentage of requests to the server serviced by the database cache, file cache, and entry cache; Availability and responsiveness of the server; Metrics revealing load on the server, such as the current number of connections to the server, the number of connections handled by the server, the operations completed and outstanding operations on the server;
	eDirectory Server	<ul style="list-style-type: none"> Monitors the block and entry caches of eDirectory with respect to database entry current size, database entry items cached, database block items cached, entry cache hit ratio etc. Provides statistics on the accesses, operations and errors of each application protocol interface of a directory server with respect to unauthorized/invalid requests, read requests, add entry requests, remove entry requests, no of chaining, no of error requests, Replication updates In and Out, incoming and outgoing traffic etc.
Authentication/Policy Servers	SiteMinder Policy Server	<ul style="list-style-type: none"> Key statistics indicating the availability of the server, and the total time taken to perform authorization and authentication checks; Statistics related to the administrator logins to the server, such as the number and percent of admin rejects, etc.; User authentication-related metrics, which include the number of authentication attempts, accepts, rejects, etc.; Performance metrics associated with authorizations, which include the number of authorization accepts, rejects, etc.

Authentication/Policy Servers	SiteMinder OneView Monitor	<ul style="list-style-type: none"> • Critical statistics pertaining to every step of the request authorization cycle, which include the number of login attempts and failed logins, the number of validation attempts and failures, the number of authorization attempts and failures, etc.; • The SiteMinder Agent Cache-related metrics, such as the percentage of resource cache and session cache hits; • Key performance data pertaining to the SiteMinder Policy server, which include the number of open connections on the server, number of authentication and authorization requests accepted and rejected by the server, etc.
	Radius	<ul style="list-style-type: none"> • External measures of availability and response time; • Monitoring of critical Radius processes
	Citrix Secure Gateway (CSG)	<ul style="list-style-type: none"> • Statistics related to connections between ICA client and the CSG, such as the number of active HTTP and ICA connections, failed connections, etc.; • Key statistics that measure to and fro traffic between the ICA client and the CSG, which includes rate of data received and sent by the CSG; • Metrics that measure validations performed by the Secure Ticketing Authority such as the rate of failed validations and the number of successful validations
	Microsoft Radius Servers	<ul style="list-style-type: none"> • Measures such as the rate of accounting requests received and responded to, the rate of duplicate accounting requests, etc., revealing how well the server performs accounting; • Key metrics indicating the health of the authorization and authentication functions performed by the server, such as, the rate at which access requests were accepted, challenged, and rejected, the rate of duplicate access requests were received by the server, etc.;
	Microsoft DFS	<ul style="list-style-type: none"> • Key metrics revealing the number of referral requests to the namespace, number of requests processed, requests that failed, request processing rate etc; • Statistics revealing the bandwidth saved during replication of each folder, number of files that were moved to the conflict and deleted folder, number of conflict files and deleted files, number of staging files that were generated, number of staging files cleaned up, deleted space in use etc; • Metrics revealing the length of the API queue, the throughput and bandwidth savings of each replication connection, the number and type of API requests; request processing rate for each type, database lookups and commits performed by each volume hosting the replication folders etc;
	Microsoft RAS Servers	<ul style="list-style-type: none"> • Port-specific performance metrics, which include, the rate at which bytes/frames were transmitted and received, the rate of errors, the number of remote access connections, etc.; • Performance statistics related to the telephone-communication on a computer, which include, the number of telephone lines serviced by the computer, the rate of outgoing/incoming calls, etc.

	Citrix Secure Ticketing Authority (STA)	<ul style="list-style-type: none"> Important metrics indicating the status of tickets such as the rate of receipt of validated and failed data requests, the rate of receipt of validated and failed ticket requests, the rate of ticket timeouts, etc.
Infrastructure Services	DNS	<ul style="list-style-type: none"> External measure of availability of the server and response time
	Windows DNS	<ul style="list-style-type: none"> External measure of availability of the server and response time; Important measures specific to a DNS server running on windows, such as the rate of requests and responses to queries, the rate of recursive query successes and failures, etc.
	FTP	<ul style="list-style-type: none"> External measures of availability and response time for GET and PUT operations
Network Devices	Network switch/router (any MIB-II compliant device) – including 3Com/Cisco/HP/BayStack switches/hubs	<ul style="list-style-type: none"> Network connectivity and response time monitoring, bandwidth usage on all incoming and outgoing links of the network switch/router
	Cisco ASA	<ul style="list-style-type: none"> CPU usage measures revealing the percentage of time during the last 5 seconds, 1 minute and 5 minute the device was using the CPU; Statistics reporting total amount of memory available for the memory pool, number of bytes from the pool that are currently used by the applications on the managed device and percentage of unused memory in the pool; Key metric revealing number of sessions that are currently active on the Cisco ASA device and number of users who have active remote access sessions; Measures indicating number of IPsec Phase-1 and Phase-2 IKE Tunnels that are currently active, number of packets received and sent by all IPsec Phase-1 and Phase-2 IKE tunnels and number of packets that were dropped by all IPsec Phase-1 and Phase-2 IKE tunnels while sending and receiving data;
	Cisco Routers	<ul style="list-style-type: none"> Network connectivity and latency monitoring; Bandwidth usage on all incoming and outgoing links; tracking of CRC errors, collisions, load factor, reliability of each of the interfaces of the router; CPU, memory, buffer utilization on the router; Hardware monitoring including temperature, voltage, etc.
	Cisco Catalyst Switch	<ul style="list-style-type: none"> Performance Statistics pertaining to a Cisco Catalyst switch, such as the rate at which data is transmitted and received by the corresponding interface of the switch, and the percentage of bandwidth utilized by that interface
	Local Director	<ul style="list-style-type: none"> Monitoring of virtual and real server workloads in terms of connection and data rates

Network Devices	Network Printers	<ul style="list-style-type: none"> • Toner capacity, paper availability, toner availability, paper flow, the number of pages printed, and printer availability; • Measures indicating the status of covers/doors, input trays and output bins
	Check Point Firewall-1	<ul style="list-style-type: none"> • Overall availability and responsiveness of the firewall; • Usage statistics including packet processing rate, percentage of rejected and dropped packets, etc.
	Cisco VPN Concentrator	<ul style="list-style-type: none"> • Fan-related metrics such as the speed and status of the individual fans; • Statistics monitoring the temperature of different hardware components such as temperature of the CPU and the cage, etc.; • Measures indicating the level of CPU, session, and throughput utilization
	HP Router	<ul style="list-style-type: none"> • Key metrics revealing the CPU and memory utilization; • Metrics revealing the configured power and the actual power of the target HP Router; • Tunnel related statistics such as the number of active tunnels, amount of data transmitted and received through the tunnels, number of packets transmitted and received through the tunnels etc; • The current voltage on each voltage test point of the managed router;
	Cisco Nexus Switch	<ul style="list-style-type: none"> • The percentage of time during the last minute/during the last 5 minutes the device was using the CPU; • The current state of each fan sensor and the speed of each fan; • The memory utilization of each memory module; • The current state of each sensor of the power supply units; • The current state of each voltage sensor; • The size of the RAM and NVRAM; • The utilization percentage of NVRAM; • The current operational status of each fan; • The availability of all network interfaces; • Metrics revealing the data transmission to and from each network interface and also error -related statistics of each network interface;
	Palo Alto Firewall	<ul style="list-style-type: none"> • The availability of the firewall and the network connectivity of the firewall; • The high availability status of the firewall and the mode of high availability configuration; • Sessions related measures such as total number of sessions that are currently active and the number of active TCP/UDP/SSL Proxy sessions; • Key metrics revealing number of tunnels that are active when GlobalProtect Subscription is enabled and the utilization of the GlobalProtect gateways; • Measures indicating the number of sessions that are active on each virtual system and the percentage of sessions that are utilized on each virtual system;

Network Devices	Blue Coat ProxySG	<ul style="list-style-type: none"> Metrics indicating the total number of client HTTP connections made by the Blue Coat ProxySG appliance, the number of active client and server connections, and the number of idle client and server connections; Measures revealing the rate at which the HTTP requests were received from the client, the rate at which the number of HTTP hits were processed by the client, the number of HTTP errors, the service time taken by the appliance to process the HTTP hits, the HTTP partial hits and HTTP misses, etc; Key metrics providing the details on the total number of unencrypted and encrypted ICAP scanning transactions performed between the appliance and the ICAP server, the number of encrypted and unencrypted requests transferred between the appliance and the ICAP server, the number of successful and failed ICAP transactions, etc; Measures indicating the resource availability status and the resource utilization of the appliance; Statistics revealing the CPU utilization of the appliance, the total amount of memory allocated to the appliance, the amount of memory utilization of the appliance, the amount of memory available for use and the percentage of memory utilization of the appliance; Temperature sensor-related metrics such as the current status and temperature of each sensor; Voltage sensor-related metrics such as the current status and voltage reading of each sensor; Fan sensor-related metrics such as the current status and speed of each sensor; Measure reveals the current operational status of each disk available in the appliance;
	Dell Switch M-Series	<ul style="list-style-type: none"> Metrics indicating CPU and Memory utilization of the stack units in the switch; Hardware related measures such as the current operational status of the fan and the current state of the power supply, and the current temperature reading of the stack unit; Measures revealing the current switch status of the stack unit; Key metric revealing the port related details such as the administrative state and operational state of each port on the switch, and strength of the power signals received and transmitted through each port.
	Cisco PIX Firewall	<ul style="list-style-type: none"> Buffer usage metrics such as the maximum blocks that were allocated, the number currently available etc.; Measures indicating status of a hardware unit; Connection-related metrics, which includes the number of open, closing, and half-open connections

Network Devices	Juniper SA Device	<ul style="list-style-type: none"> Key host-level statistics such as the percentage of disk space, CPU, memory, and swap space utilized on the IVE system; Critical performance metrics such as a measure of the log file growth on the IVE system, the number of concurrent users on the system, the number of file/web/applet/terminal hits on the system, etc.
	3COM Core Builder Switch	<ul style="list-style-type: none"> Key metrics pertaining to each of the network interfaces supported by the switch, such as, availability of the interface, the traffic handled by the interface, its speed, etc.
	Juniper DX Device	<ul style="list-style-type: none"> Key host-level statistics such as the percentage of disk space, CPU, memory, and swap space utilized on the IVE system; Critical performance metrics such as a measure of the log file growth on the IVE system, the number of concurrent users on the system, the number of file/web/applet/terminal hits on the system, etc. Metrics reporting the level of HTTP traffic on the device, and HTTP errors encountered by the device
	Juniper EX Switch	<ul style="list-style-type: none"> Metrics revealing the CPU utilization, current temperature, uptime of each hardware component; Memory related metrics such as the total memory allocated to each hardware component, buffer memory and heap memory utilization of each hardware component; Metrics revealing the number of times the temperature failure, fan failure and power supply failure events were triggered; The current mode of the routing engine can be ascertained;
	OpenVPN Access Servers	<ul style="list-style-type: none"> Measures revealing the total number of active users connected the server and the user related statistics such as the data transmission and reception, active session duration of each user etc; Statistics revealing the total number of licenses available for the server, the number of licenses used and the number of licenses available for use;
	QNAP NAS Systems	<ul style="list-style-type: none"> Disk related measures such as the current state of the disk, capacity of the disk and the current temperature of the disk; Metrics revealing the current state of each disk volume and the space utilization metrics of each disk volume such as the total volume, volume available for use, volume already used etc; Current speed of the fan; System related metrics such as the current CPU utilization, temperature etc; memory related metrics of the system such as the total memory, available memory, etc; The Uptime and reboot details of the system;

Network Devices	Data Domain	<ul style="list-style-type: none"> Hardware related measures such as the status of each fan, fan speed, status of the disks, disk capacity and measures related to disk I/O, size of the NVRAM and errors encountered by the NVRAM, battery status etc, Metrics revealing the current status of each Power module and sensors, overall CPU utilization and disk I/O of the storage system etc, Key metrics revealing the current status of the file system service, the space utilization of each file in the storage system, size of each file prior to and after compression, compression ratio of each file etc,
	Big-IP/F5 Load Balancer	<ul style="list-style-type: none"> Measures key system performance metrics for any system that supports the Net-SNMP agent, and checks for the availability and used memory of a system using the Net-SNMP agent; Reports the status of incoming and outgoing traffic through all the virtual servers/addresses that have been configured on the load balancer; Reports the status of and the traffic on each of the virtual IP addresses configured on the BIG-IP load balancer; Monitors the status of and the traffic on each of the virtual servers configured on the BIG-IP load balancer
	Brocade SAN Switches	<ul style="list-style-type: none"> Metrics indicating the current status (whether online or offline) of the fabric switch, and the status (whether online/offline/faulty/unknown/testing) of the links of the ports on the switch; Key statistics revealing the health of sensors, such as, the number of sensors in the normal/unknown/faulty state; Metrics indicating the status of the ports on the fabric switch;
	FortiGate Firewall	<ul style="list-style-type: none"> Usage statistics indicating how well the firewall uses the CPU, memory, and disk resources; Statistics measuring the load on the firewall in terms of the number of sessions to the firewall, the network and data traffic processed by the firewall; Firewall efficiency indicators such as the number of attacks and viruses that were detected by the firewall during the last 20 hours
	Alcatel Switch	<ul style="list-style-type: none"> Overall health statistics pertaining to the switch such as, input utilization, I/O utilization, CPU, memory utilization, etc.; Usage statistics of each of the modules of the switch such as, input utilization, I/O utilization, CPU, memory utilization, etc.; Metrics revealing the status and usage of ports on the switch.

Network Devices	Forefront TMG	<ul style="list-style-type: none"> Measures relating to the content cache of the firewall such as the rate at which the data is retrieved from the disk allocated for content caching; the rate at which the I/O operations failed, how well the URLs are stored and retrieved from the disk drive; etc., Metrics revealing the efficiency of the firewall such as the number of scanned messages, the number of infected messages that were blocked, the number of spam messages etc., Metrics revealing the number of packets that were backlogged, dropped, allowed etc., by the firewall, etc., Metrics revealing the number of connection objects that were waiting for a TCP connection, number of active sessions, SIP registrations, number of active SIP sessions, TCP and UDP connections, the data transmission and reception rate through the firewall, number of worker threads etc., Measures indicating the number of active H.323 calls and the rate at which the calls were handled; Measures revealing the number of connections that were active, connecting, listening through the SOCKS protocol, Metrics revealing the data transfer between the web proxy clients and servers such as the number of HTTP/HTTPS requests, rate at which data is sent from the computers protected by the firewall, number of outgoing and incoming connections etc.,
	Citrix Branch Repeater	<ul style="list-style-type: none"> Measures revealing the uptime and the percentage of CPU used by the branch repeater; Metrics revealing the data transmission and packet transmission for each application link, the numerical statistics of the dropped data and packets during transmission as well as reception due to QoS threshold settings; Measures indicating the load on the branch repeater, the effectiveness of the repeater is reported by the amount of data and connections it has accelerated, the compression rate of the accelerated date etc., Measures indicating the number of active connections, metrics revealing the accelerated traffic to and from each ICA application, the ratio of data transmission and reception for each ICA application etc., Metrics revealing the level of traffic sent and received by each configured service class, the packets dropped when QoS thresholds are violated by each service class, etc., Metrics revealing the volume of traffic handled by each WAN and LAN link and the packets dropped over the links due to QoS threshold violation etc., Measures revealing the traffic accelerated by each user-configured traffic – shaping policy and the rate of traffic acceleration and the packet loss due to each traffic shaping policy etc.,
	Cisco SAN Switch	<ul style="list-style-type: none"> The status of the fan tray, the power supply units, and the sensors; Metrics such as the operational mode and operational state of each WWN of the fibre channel port, the link failures experienced by each WWN, etc.,; The operational status of each VSAN on the fibre channel switch; Overall health statistics pertaining to the switch such as, input utilization, I/O utilization, CPU, memory utilization, etc.;

Network Devices	Cisco CSS	<ul style="list-style-type: none"> • Metrics revealing the current status and usage of the destination services configured for every service group, which include, the current destination service status, the number of time user requests to the service group load balanced to the service, etc.; • Metrics revealing the status and usage of each service group, such as, the current status of the groups, the number of connections established through the groups, etc. • Statistics related to the load generated by applications connecting to the CSS, such as, the number of application packets received and transmitted, the current session state, etc.; • Measures that reveal owner activity on the device, such as, the number of times the owner accessed the CSS with content requests, the number of owner requests dropped, etc; • Key metrics pertaining to the status and usage of the content providing services, such as, the current state of each service, the total data sent through each service, the average service load, the service status, etc.; • Service-related statistics such as, the number of times each service served owner requests, the status of each content service, etc.; • Metrics revealing the state and usage of content rules; • Metrics indicating the current state of each IP interface on the CSS, and the types of interfaces associated with each VLAN circuit configured on CSS; • Overall health statistics pertaining to the switch such as, input utilization, I/O utilization, CPU, memory utilization, etc.; • The current status of each redundant link on the CSS.
	Coyote Point Equalizer	<ul style="list-style-type: none"> • External metrics of availability and responsiveness of the equalizer; • Cluster-related statistics, such as, the load on each cluster, the total number of connections handled by each cluster, the rate at which servers in each cluster were accessed, etc.; • Metrics revealing the number and nature of connections to the equalizer; • Statistics related to cluster servers, such as, the load on each server in each cluster, the responsiveness of each cluster server, the idle time of each server cluster, etc.
	Generic SAN Switch	<ul style="list-style-type: none"> • The current status of the sensors on the switch, the connection units on the switch, etc.; • Port-related statistics, such as the load on each port on the SAN switch, the current operational state and health of each port, the port speed, the errors captured on every port, etc.,

Network Devices	Coyote Load Balancer	<ul style="list-style-type: none"> Metrics revealing the total connections, active connections, total transactions of each server, amount of data received/transmitted through the server, number of HTTP responses compressed by the server; Metrics revealing the current state, failover state and failover mode of each peer; Measures indicating the total connections, active connections, total transactions of each server, amount of data received/transmitted through each HTTP cluster and HTTPS cluster, number of HTTP responses compressed by the HTTP cluster and HTTPS cluster, etc; Measures indicating the total connections, active connections, total transactions of each server, amount of data received/transmitted through the L4 cluster; Measures indicating the total connections, active connections, total transactions of each server, amount of data received/transmitted through the L4 cluster; Measures indicating the total connections, active connections, total transactions of each server, amount of data received/transmitted through each server pool; Metrics revealing the current state and subnet state of each VLAN;
	F5 Big-IP Local Traffic Manager	<ul style="list-style-type: none"> Metrics indicating the temperature and fan speed of the CPU supported by the local traffic manager; Metrics indicating how the traffic manager utilizes each disk partition; Chassis fan-related metrics such as the current status and speed of the fan; The current status of each pool and virtual server configured on the traffic manager The connections, data and packet transmission/reception through each virtual server; The maximum number of connections that were established on each virtual server; The availability of each node; The activity status of each load balancing pool if the pool is available; The connections, data and packet transmission/reception through each pool and pool member; The maximum number of connections that were established on each pool and pool member; The data and packet transmission/reception from the client and server on the Traffic Management Module; The client and server related connections on the Traffic Management Module etc., The connections, data and packet transmission/reception through each virtual server; The maximum number of connections that were established on each virtual server;

Network Devices	Sonic Firewall	<ul style="list-style-type: none"> Measures indicating the active connections on the firewall and the maximum connections that can be handled by the firewall; Measures revealing the reception/transmission of fragmented packets through each VPN tunnel, amount of data and packets encrypted/decrypted by each VPN tunnel; Measures revealing the current CPU utilization and memory utilization of the firewall;
	WatchGuard Firewall	<ul style="list-style-type: none"> Measure revealing the current CPU utilization of the firewall at periodic intervals; Measures indicating the total connections, active connections and the connections that were dropped by the firewall; Measures revealing the data and packet transmission/reception through the firewall; For each firewall policy, measures such as the connections, data transmission/reception, packet traffic and packets discarded due to errors is reported; For each VPN pair, measures such as the IPSEC traffic, packet traffic, packets discarded due to errors is reported; Data traffic, packet traffic, packets discarded due to errors etc., are reported for each VPN tunnel; Measures revealing the Data traffic through various protocols such as Encapsulating Security Payload (ESP), Authentication Header, IP Payload Compression protocol (IpComp) of each security association; Measures revealing various errors during packet transmission for each security association etc.,
	QNAP NAS	<ul style="list-style-type: none"> Measures indicating the current status, capacity and temperature of each disk of the QNAP NAS system; Measures indicating the current status of each disk volume and measures relating to its utilization; Measures revealing the speed of each fan available in the QNAP NAS system, the CPU and memory utilization of the QNAP NAS system, the uptime of the system etc.,

Network Devices	A10 Application Delivery Controller	<ul style="list-style-type: none"> Measure revealing the percentage of CPU resources utilized; the current state of each power supply unit; Measures indicating the total capacity of the disks, amount and percentage of space utilized from the total capacity of the disks and free space available in the disks; Memory related metrics such as total amount of the memory configured, amount and percentage of memory space and the amount of memory space that is available for use; For each real server/virtual server, metrics revealing the current health, rate at which data and packets were transmitted from and received by each real server/virtual server, number of active connections, percentage of connection used by each real server, number of persistent connections, L7 requests processing rate etc.; For each port of the real server/virtual server, key metrics revealing the current, the rate at which data and packets were transmitted/received through each port, number of active connections, maximum number of connections that were established through each port of the real server, number of persistent connections etc.; For each service group/service group member, key metrics revealing the current health/state, rate at which data and packets were transmitted /received, percentage of active connections that were established, the rate at which the connections were established, maximum number of connections that were established through each port of the virtual server, number of persistent connections, etc.
	CheckPoint Smart Appliance	<ul style="list-style-type: none"> Disk measures such as total space allocated for each disk, space that is currently available for use, and amount and percentage of space that is currently in use. Measures indicating speed at which each fan operates and whether/not the sensor of the fan is out of range. Metrics revealing the current status of each power supply unit, the current voltage of each element and whether/not the LED corresponding to the element is out of range, the current temperature of the hardware unit. etc CPU related metrics such as percentage of CPU utilization and CPU utilized for system level and user level processing. Memory related measures such as total memory of the appliance, amount and percentage of memory utilized by the appliance and amount of memory that is currently available for use in the appliance. Measures indicating percentage of CPU utilized by each virtual system, the number of connections that are currently active on the virtual system and the maximum number of connections made to the virtual system.

Messaging Servers	MSMQ	<ul style="list-style-type: none"> Monitoring server performance through key metrics like incoming and outgoing message rates, error messages seen by the server, current sessions being handled, etc. Queue statistics, such as, the outstanding message queue length, and the journal queue length
	WebSphere MQ Series	<ul style="list-style-type: none"> Monitoring the availability of the messaging service and the individual queue managers; Statistics pertaining to the local queues including the current depth, percentage utilization of the queue, details of which messages are in the queue and for how long, etc.; Measures that indicate channel availability and volume of data handled by each channel
	Fiorano MQ	<ul style="list-style-type: none"> Server statistics such as the number of active and idle threads, the rate of receipt of messages from and publication of messages to topics, the rate at which messages are received from and pushed to the queues on the server, memory statistics including the amount of memory utilized and the available memory; Queue-specific statistics such as the number of messages in the queue awaiting delivery, and the number of undeleted messages in the queue; Topic-related measures such as the number of durable subscribers on the topic, and the number of undelivered messages to such subscribers
	Microsoft Skype for Business / Lync	<ul style="list-style-type: none"> Key metrics that reveal the address book accesses, the speed of these accesses etc., statistics revealing the message processing ability of the server, Metrics that capture the failed messages and when the failure occurred – during message validation? in the message queue or when written to the database? Metrics revealing the quality of the AV Conferencing experience, metrics revealing the processing of the call park requests; Statistics revealing the number of users/clients connected to the server, the current load on the server, the number of dropped conferences and unfinished tasks; Metrics revealing the number of requests queued for processing by the database and the time spent by the requests in the queue; Key metrics that detect the client and server authentication failures as well as DNS failures; The current state and draining state of the Application Sharing Conference unit, the Data MCU unit and the Instant Messaging Multipoint Control unit;

Messaging Servers		<ul style="list-style-type: none"> Key metrics revealing the number of RDP connections that failed and the conferences handled by the Application Sharing Conference unit; The number of whiteboards and conferences that were active on the Data MCU; The number of add user and add conference responses that failed; the number of SIP connections and the that failed and the connections that were currently active; Key metrics revealing the number of replication requests received by the replication service and the time taken by the service to process the requests; The number of times user provisioning and publish calls failed; Key metrics revealing the number of stored procedure calls executed by the user service; the number of unique users connected to the servers based on each client version; Statistics revealing the users for whom voice call is enabled, the users who are currently connected to the server; the number of times the HTTPS connections failed etc.;
	Tibco EMS Server	<ul style="list-style-type: none"> Metrics that indicate the current mode of operation of the server; Metrics revealing the load on the server in terms of the number and type of topics, and the number and type of queues on the server; Metrics that report the percentage of connections to the server that are in use, and the number of producers, consumers, sessions, and durables on the server; Metrics revealing the number and size of pending messages on the server, the free and used message memory, etc. Metrics related to each queue on the server, such as, the number and size of pending messages to a queue, whether the queue is a static or dynamic queue, etc. Metrics related to each topic on the server, such as, the number and size of pending messages to a topic, whether the topic is a static or dynamic topic, etc. Metrics that reveal the level of user activity on the server, such as, the number of sessions and connections to a user, etc.
Mail Servers	Microsoft Exchange 2000/2003	<ul style="list-style-type: none"> Extensive server statistics including the number of logons that have failed, connection and transmission rates, authentication failure rates, etc.; Key message delivery statistics such as the number of messages in the local and remote queues, current number of inbound SMTP connections, number of messages entering the retry queue, average number of retries per outbound message, local and remote message delivery times, etc.; MTA-related statistics, such as, the message processing rate of the MTA, the number of threads in use by the MTA, the number of outstanding messages in the work queue, etc.;

		<ul style="list-style-type: none"> Metrics related to the mail scanning and virus processing sub-system, which includes, the number of requests pending virus scanning, the rate at which messages were scanned, the number of files and messages that were cleaned and quarantined, etc.; Measures monitoring the performance of RPC activities, such as, the number of attempted RPC calls, the number of RPC failures and successes, the RPC success ratio, RPC latency-related metrics, etc.; Mailbox usage metrics such as, the mailbox size, quota size, and percentage of mailbox space currently used.
Mail Servers	Microsoft Exchange 5.5	<ul style="list-style-type: none"> Caching-related statistics such as the number of requests fulfilled by the Exchange Directory cache; Database usage statistics, which include the number of requests fulfilled by the buffer pool, the percentage of tables opened using the cached information, etc.; Metrics pertaining to the POP3 and IMAP protocols that include the number of current IMAP and POP3 connections, and respective outstanding requests; General statistics including the number of incoming and outgoing messages from the mail server, the number of messages awaiting final delivery, etc.
	Exchange 2007/2010	<ul style="list-style-type: none"> External metrics revealing the availability of the server to send and receive mails, and its responsiveness; Metrics indicating the overall request processing ability of the server, such as, the rate at which the server processes requests, and the speed at which requests are processed; Metrics revealing the health of the ActiveSync engine, such as, the time taken by the engine to process requests, the number of requests in queue, the number of pending ping commands, the number of busy/idle worker threads, etc.; Metrics revealing the health of the Availability service, such as, the request servicing rate, the rate of mailbox hits/misses, etc.; Metrics related to the Outlook Web Access protocol, such as, the speed of the web access, the time taken for a search to complete, the number of failed web access requests, etc.; Metrics monitoring the Exchange server-AD server interactions, such as, the rate of AD cache hits and misses, outstanding LDAP requests to the AD server, the rate at which LDAP searches were timed out, the rate of LDAP fatal errors, etc.; Metrics related to the RPC activity on the server, such as, the rate of RPC successes and failures, the RPC success ratio, the rate of high latency RPC attempts, the fast RPC ratio, etc.; Metrics monitoring the health of the Exchange database, such as, the database cache hit and table cache hit ratios, the rate of log records waiting to be added to the log buffer, and the rate of log threads waiting for their data to be written to the log buffer;

Mail Servers	<ul style="list-style-type: none"> • Metrics indicate the mount status of the Exchange mailbox; • Key statistics pertaining to the health of the mailbox store and public store, such as, the time taken by the mailbox store to deliver a message to local recipients, the count of messages in queue, the rate of messages delivered, the size of the receive queue, etc.; • Usage metrics related to the virtual memory of the Exchange store, such as, the size of the largest free block of virtual memory, the number of free blocks of virtual memory, etc.; • Metrics related to the Search Indexing engine, such as, the number of mailboxes left to crawl, the number of documents to be indexed, the number of documents that failed during indexing, the time taken for document indexing, etc.; • Metrics that shed light on issues related to the Mailbox assistant, such as, the number of events awaiting processing by the assistants, delay between polling and event creation, the time taken by the assistant for event processing, etc.; • Metrics related to the mailboxes in a storage group, which include, the number of users who are currently reaching their mailbox quota, the number of users who have reached their storage quota, etc.; • Queue-related statistics, such as, the number of messages currently in queue; • Metrics measuring the effectiveness of the Recipient Filter Agent, such as, the number of recipients rejected by recipient validation and by block list; • Statistics related to the Sender ID agent, such as, the rate at which messages of various types are validated; • Metrics that monitor the queues on the server to reveal its processing ability, such as, the number of messages in the active, retry, and remote delivery queues, etc.; • Metrics indicating the health of the SMTP Receive Connectors, such as, the number of messages received by the server, the messages refused owing to size, etc.; • Metrics indicating the health of the store drivers, such as, the number of failed and successful deliveries, the rate of inbound message delivery attempts, etc.; • Metrics tracking the connection filtering agent's activities, such as, the number of connections to IP block list providers, IP allow list providers, IP block list, and IP allow list; • Metrics that reveal the level of efficiency of the content filter agent's operations, such as, the number of messages at spam control level 0-9, the number of messages quarantined, the number of messages deleted, etc.; • Metrics that report the number of senders blocked by the Sender reputation agent owing to various reasons; • Metrics that reveal the level of traffic handled by the SMTP send connectors, such as, the rate at which messages were sent by the connectors, data transferred per connection, etc.;
---------------------	---

	Microsoft Exchange 2013/2016	<ul style="list-style-type: none"> • Metrics indicating the LDAP read rate, LDAP read time, LDAP search rate and time, timeout errors. Outstanding requests etc., for each client process communicating with the Active Directory server; • Metrics indicating the LDAP read rate, LDAP read time, LDAP search rate and time, timeout errors. Outstanding requests etc., for each domain controller communicating with the Active Directory server; • Key metrics indicating the authentication requests, latencies and rejections for every configured authentication method; • Metrics revealing the count of the items in the mailbox, size of the mailbox, quarantined status of the mailbox, various metrics pertaining to the quota configuration of the mailbox, accessibility of the mailbox etc., • Key metrics pertaining to the databases of the exchange server such as the database size, mount status, mailboxes in the database etc can be ascertained; • Metrics indicating the quality check on the indexing process such as the count of documents that skipped indexing, documents that failed indexing etc.; • Metrics related to the Search Indexing engine, such as, the number of mailboxes left to crawl, the number of documents to be indexed, the number of documents that failed during indexing, the time taken for document indexing, etc.; • Metrics monitoring the mailbox databases and reports the type of replication set, health of the database copies, log files that are pending to be copied, disk capacity of the database etc, • Metrics related to the current status of the mail flow, time taken for mail delivery, key metrics revealing the user experience with POP3 service such as the number of SSL connections, command processing rate, LDAP latency POP3 connections that are currently open etc; • Key metrics revealing the replication health of the Exchange server, current status of the Exchange search on the mailbox, search time, etc; • Statistics relating to the ESE databases such as the database cache size, database cache hit ratio, database read and write latencies, page compression ratio in the cache of the database, database cache size, read/write operation latency in the cache, database sessions in use etc; • Key statistics pertaining to the health of the exchange store such as the load on the database, the processing ability of the database, overall health of the database, the number of mailboxes that are quarantined in the database, number of active mailboxes, database maintenance rate, mailbox maintenance rate, message property promotion rate of the database, etc; • Key metrics revealing the client types/protocols communicating with the Exchange store such as the RPC request load, operational load, log file load processing time, growth, etc,
--	-------------------------------------	---

Mail Servers	<ul style="list-style-type: none"> • Key metrics pertaining to the anti malware scan engines such as the errors in the engine, number of malware items, scan time per item etc, • Statistics revealing the maximum time a component took to process email messages, key metrics pertaining to the delivery SMTP receive connectors/SMTP send connectors such as the incoming load on the connector, the rate at which each connector processed the load, and the count of mails rejected by the connector etc, • Metrics that shed light on issues related to the Mailbox assistant, such as, the number of events awaiting processing by the assistants, delay between polling and event creation, the time taken by the assistant for event processing, .etc; • Metrics related to the submissions made by the mailbox transport submission service such as the failed rate, successful submissions rate, permanent failed submissions, failures encountered by store driver submission agents, the time taken by the submission component to process messages for various latencies etc, • Statistics revealing the maximum time a component took to process email messages, key metrics pertaining to the submission SMTP send connectors such as the throughput of the connector, errors encountered by the connector; • Statistics revealing the maximum time a component took to process email messages, key metrics pertaining to the submission SMTP send connectors such as the throughput of the connector, errors encountered by the connector; • Metrics revealing the errors encountered by the classification scan engine, time taken by the engine to load and scan the content etc, • Metrics revealing the maximum time taken for end to end email message flow for various latencies, number of messages of each priority in various exchange transport queues, the processing rate of messages of every priority in the queue etc, • Statistics pertaining to the current state of the queue, messages in the queue, rate at which messages are sent/received etc, • Statistics pertaining to the number and size of messages pertaining to each key event type handled by the exchange server, number of messages deferred due to errors in the transport rule evaluation, the maximum time taken by each transport component to process messages for various latencies etc, • Key metrics indicating the load on the SMTP receive connector/SMTP send connector, messages rejected by the connector, errors encountered by the send connectors etc, • Key metrics indicating the time taken by the Client Access server to process the requests to each proxy enabled service, the length of the HTTP service request queues, rate at which the requests are rejected, the cache utilization of the proxy enabled services etc, • Metrics revealing the maximum time taken by the Frontend transport connector took to proxy emails at various latencies, key metrics indicating the load on the Frontend SMTP receive connector/SMTP send connector, messages rejected by the connector, errors encountered by the send connectors etc, • Metrics revealing the size of the mailbox folder and its sub folders, number of items in the folder and its sub folders etc, • Statistics revealing the availability of the Exchange mail server and performance of the servers, metrics tracking the synchronization and ping requests to the Activesync servers, the incoming and outgoing proxy requests through the Activesync servers etc,
---------------------	---

Mail Servers		<ul style="list-style-type: none"> • Statistics revealing the requests for Outlook WebApp Services and instant messaging services, time taken to service such requests, number of requests that failed, unique users in the Outlook WebApp etc, • Key statistics pertaining to the RPC connection attempts to the server over HTTP, connection failures, RPC packet rate etc, • Metrics revealing the number of users currently connected to the Exchange server over RPC/HTTP, incoming/outgoing bandwidth of the RPC/HTTP requests, number of connections made to the Exchange server instance through the RPC/HTTP etc, • Statistics revealing the inbound calls to the Unified Messaging Call router service, number of inbound calls received and rejected, the overall health and performance of the Unified Messaging Call router service;
	Instant Messenger of the Exchange 2000 Server	<ul style="list-style-type: none"> • Performance metrics related to messenger such as the count of users currently online, the number of messages in queue, the rate of request failure, the rate of response failure, the requests receive rate, the requests sent rate, etc.
	SunONE Messaging Server	<ul style="list-style-type: none"> • Performance statistics related to the server's HTTP, IMAP, and POP3 services, such as the number of active connections and sessions to the server, the rate of successful logins and failed ones, etc.; • External measures indicating the availability and responsiveness of the TCP ports of the messaging server's LDAP, IMAP, and POP3 services; • Measures of the database lock behavior, such as the rate of lock requests, releases, and deadlocks, and transaction metrics such as the rate of transaction begins, commits, rollbacks, etc.; • Message Traffic related metrics such as the number of messages delivered, rejected, etc.; • User-specific statistics such as the number of user accounts present, the disk space consumption of the messages in the user accounts, etc.
	Lotus Domino	<ul style="list-style-type: none"> • Statistics pertaining to the Domino database such as number of databases in cache, the percentage of pages cached, the maximum capacity of the server cache; • Memory usage metrics, which include the total memory allocation, the shared memory allocation, etc.; • Statistics related to the services of the Domino mail server such as the message receipt rate, percentage of undelivered messages, etc.; • Important metrics pertaining to the network traffic, that includes, the data receipt and transfer rates, etc.

Mail Servers	Novell Groupwise	<ul style="list-style-type: none"> • Tracking the availability and responsiveness of all of the Groupwise server components including the Groupwise Internet Agent, Message Transfer Agent (MTA), Post Office Agent (POA), and Web Access Agent; • Performance metrics such as rate of data/messages sent and received, number of messages in each of the server queues (routing queue, post office queue, gateway queue), etc.; • SMTP service related metrics including the number of send and receive threads available, rate of message transmission and reception, etc.; • POP3, LDAP, and IMAP service related measures such as the number of sessions active, etc.; • Client connection-related measures that include request rate, pending request count for POP3 access; • Local queue related metrics such as the number of messages in the routing queue, post office and gateway queues, etc.
	Qmail	<ul style="list-style-type: none"> • Statistics revealing the availability and down time (if any) of the services on a Qmail MTA; • Queue-related metrics such as the total number of messages in the queue, the queue size, etc.; • Key metrics pertaining to the delivery performance of the server such as the number and rate of delivery attempts that were made, the number and rate that was successful, etc.; • Message-related statistics, which include the number of messages transferred, bounced, thrown away etc.
	Ironport AsyncOS Mail	<ul style="list-style-type: none"> • Overall health statistics revealing the CPU utilization, Disk I/O utilization, Power supply status, fan speed and current temperature of the server; • Queue-related metrics such as the used percent of total queue capacity and the available queue space details. • Key metrics pertaining to the DNS statistics of the server such as the outstanding DNS requests and the pending DNS requests; • Mail and socket related metrics, which include the open socket count and mail thread count.

Mail Servers	BlackBerry Server	<ul style="list-style-type: none"> Metrics related to the BlackBerry MDS Services component of BlackBerry MDS, which includes, the number of connections to handheld devices and number of push server connections initiated by MDS, the data traffic to and from handheld devices, the number of packets refused, invalid packets, failed connections, truncated connections, successful connections, etc. Metrics revealing how well the BlackBerry server handles messages sent and received from handheld devices, such as, the number of messages processed, the number awaiting delivery, the number of messages that expired or were undelivered, etc. SRP connection-related metrics revealing whether the BlackBerry server/BlackBerry Dispatcher is connected with the SRP host or not, the number of times reconnection attempts failed, the duration for which the connection with the SRP host was lost, etc.; Statistics that indicate the license usage by users to the BlackBerry server, such as, the number of installed, used, and free licenses; Metrics that bring to light the efficiency of the BlackBerry Messaging Agent, such as, the response time for operations, number of failed connection attempts to the mail server, etc. Measures that indicate whether WER (Wireless Email Reconciliation) is enabled on the BlackBerry server or not, whether the handheld device is in the cradle or not, whether a user account is enabled or not, etc.; Measures of the user activity on the BlackBerry server, which include, the number of messages processed, the number awaiting delivery, the number of messages that expired or were undelivered, per user;
	Postfix	<ul style="list-style-type: none"> Metrics revealing the total size of the deferred queue and the number of messages that were in the queue during different time slots; Hold queue related measures such as the total size of the queue and the number of messages that were in the queue during different time slots; Messages relating to the MailDrop queues such as the total size of the queue and the number of messages that were in the queue during different time slots; Active queue related issues such as the total size of the queue and the number of messages that were in the queue during different time slots;
Microsoft Windows Applications	BizTalk	<ul style="list-style-type: none"> Key metrics at EVERY stage of server processing including document submission and reception rates, encode/decode rates, the rate at which the documents are being mapped, the rate at which the documents in the work queue are being parsed, etc.
	DHCP	<ul style="list-style-type: none"> Performance statistics pertaining to a DHCP Server, that includes, packet processing time, length of the internal message queue, rate of receiving and acknowledging requests, etc.; Utilization measures including the number of IP addresses in use and the number of free IP addresses in the target network

Microsoft Windows Applications	WINS	<ul style="list-style-type: none"> • Measures indicating the total number of queries received by the WINS server, rate of failed queries, release requests received, release failures, etc.
	Print Server	<ul style="list-style-type: none"> • Measures related to print queues such as the number of jobs, job service rate, number of errored jobs, etc. of a print queue
	Transaction Server	<ul style="list-style-type: none"> • Monitoring of CPU and memory usage of the server as well as statistics on packages installed and those that are running; • Transaction-related metrics such as number of aborted and committed transactions
	Proxy Server	<ul style="list-style-type: none"> • Measures such as the availability and response time of the server, the TCP connection availability, etc.; • Measures related to the WinSock service including the active sessions for the service, the number of live worker threads etc.; • Health of the Web Proxy service in terms of the number of active sessions to the web proxy service, the time taken by the service to service a request, etc.; • Caching-related metrics that include, the refresh rate of the URL cache, the cache size, etc.
	ISA Proxy Server	<ul style="list-style-type: none"> • Health of the firewall protection to the server in terms of the number of active TCP and UDP connections to the server, read and write rates, etc.; • Caching-related metrics that include, the commit rate of the URLs, the space utilized by the disk and memory caches, etc.; • Measures related to the Web Proxy service such as the percentage of successful client requests to the server, the time taken for processing requests, requests rejected
	Domain Controller	<ul style="list-style-type: none"> • Measures that indicate the availability and response time for domain authentications; • Various session related statistics like, rate of logons to the server, number of errored logons, the number of active, normally terminated, and timed out sessions etc.
	Event Logs	<ul style="list-style-type: none"> • Statistical information about the events generated by various applications, windows services and drivers in the system, which includes the number of application error events, application information error events, warnings, system error events, etc.
	COM+ Applications	<ul style="list-style-type: none"> • Metrics that reveal the health of the COM+ applications on the COM+ server, such as the CPU and memory usage of the application, the number of threads running in the application, etc.; • Component-specific measures, which include the number of objects being invoked, the average duration of method calls, the number of completed and failed calls, etc.; • Transaction-related metrics, which include the number of active, aborted, and committed transactions, the average response time to a user request, etc.

Microsoft Windows Applications	ASP.NET	<ul style="list-style-type: none"> • Worker process-related measures, which include the number of applications currently running, requests handled, requests rejected, request execution time, worker processes currently running, etc.; • Statistics related to the managed locks and threads used by an application, which include the number of currently managed thread objects, the number of native operating system threads, rate at which threads attempt to acquire locks, etc.; • Measures that assess the memory allocation activity of the server, such as the heap memory usage, time spent on garbage collection, etc.; • Metrics that reveal the performance of the ASP.NET application/application domain cache like the cache hit ratio, total entries in the cache, etc.; • Measures indicating how well the appdomain handles requests, such as the number of requests currently executing, the number of successful requests, requests that timed out, etc.
	Microsoft System Management Server (SMS)	<ul style="list-style-type: none"> • Metrics related to the health of the Data Discovery Manager, which include, the number of data records processed and those that are in queue; • Key measures of the health of the SMS Memory Queue, such as, the number of objects added or removed from the queue; • Software Metering-related metrics, such as, the number of software metering usage files processed by the Software Metering Processor, the count of bad usage files processed, etc.; • Key statistics that indicate the report processing ability of the hardware and software inventory managers, which include, the number of reports processed and the rate at which they were processed; • Metrics revealing how efficiently the SMS Policy manager services requests, such as, the rate at which requests are received by the GetPolicy component, and the number of requests that were served from the cache;
	File Server	<ul style="list-style-type: none"> • Key metrics that track open file connections to the host such as the number of files locked at the host, and the number of users having open files on the host; • Session-related statistics, which include the number of files opened over the network and number of users with open sessions
	Microsoft SharePoint	<ul style="list-style-type: none"> • Metrics that reveal the processing ability of the archival plugin component, such as, the number of documents that are actively using the first queue of the plugin, the number of documents actively using the second queue of the plugin, the number of documents which currently returned errors from the plugin, the number of blocked documents, etc.; • Metrics related to the document conversion process, such as, the rate of email message processing, the number of pending document conversions, etc.; • Statistics that reveal the health of the Excel calculation service, such as, the rate of requests with errors, the average processing time for a request, the Excel calculation service workbook cache size, etc.,;

		<ul style="list-style-type: none"> Measures that track the number and rate of requests to the Excel Web Services component; Metrics revealing the efficiency of the Search feature offered by SharePoint, such as, the number of queries to the content index that currently failed, the number of queries that succeeded, the number of documents filtered, etc.; Metrics revealing the effectiveness of the SharePoint publishing cache, such as, the cache hit ratio and the number of items removed from the cache; Measures that reveal the content processing ability of the gatherer, which include, the rate of document additions, the rate of documents with errors, the number of times access to a document has been retried, the number of unprocessed documents in the gatherer, etc.
Microsoft Windows Applications	Microsoft Active Dynamix	<ul style="list-style-type: none"> Metrics revealing data and session load, such as, the total number of sessions and the number of active sessions, the total number of client requests, the rate at which client requests were processed, bytes sent and received by server, etc.; Metrics related to .NET business connectors, such as, the number of .NET business connector sessions, the number and type of session-related exceptions, the number of sessions allocated, disposed, etc.
	Microsoft Dynamics NAV	<ul style="list-style-type: none"> Metrics revealing the number of client sessions to each NAV server instance, the time taken for server operations, the number of open connections to each server instance, number of rows in all temporary tables for each server instance; the percentage of requests made by each instance that were serviced by the calculated fields cache, SQL command cache, etc, Key metrics monitoring the call and transaction load on each instance of the NAV service such as the number of pending calls, failed calls, aborted transactions, etc;
	Microsoft Dynamics CRM 2011	<ul style="list-style-type: none"> Measures revealing the authentication statistics of the server such as the number of unsuccessful authentication requests per minute, number of authentication requests that were processed, number of authentication requests that were processed/failed while processed using the active directory authentication credentials, number of authentication requests that were processed/failed while processed using the authentication credentials of the Microsoft account, number of authentication requests that were processed/failed using claim based authentication etc., Critical metrics revealing the email messages passing through the Email router such as the number of corrupted incoming email messages, incoming email messages that failed during delivery, incoming email messages that were not delivered successfully, the number of times the email router service configuration was refreshed, the number of times the service provider could not be loaded/ failed to load during execution etc., Metrics revealing the total number of cache flush requests that were successfully received for the locator service, the cache flush requests that were unsuccessful etc.,

		<ul style="list-style-type: none"> • Numerical statistics dealing with the router requests such as the router requests that timed out, router requests made through faulted channels, requests received by the router, requests to the AppFabric that timed out. • Metrics related to the Sandbox host such as the total CPU percentage used by all the worker processes, memory used by all the worker processes, number of handled used by the worker processes, the organizations on which the worker processes are active, rate at which the incoming custom workflow activities are executed, SDK request related statistics such as the rate of outgoing SDK requests, percentage of SDK requests that failed, response time etc., • Critical measures related to the Microsoft Dynamics CRM Web service such as the number of requests received by the Microsoft Dynamics CRM web service, percentage of requests that failed, total number of metadata requests received by the Microsoft Dynamics CRM web service, percentage of metadata requests that failed, total number of rendering requests etc.,
Microsoft Windows Applications	Fast Search for SharePoint 2010	<ul style="list-style-type: none"> • Measures related to each crawl collection, such as, the number of websites or web links that are currently crawled, the rate at which the documents are currently downloaded, the average size of the documents downloaded, the number of documents downloaded that are currently stored in the Web Crawler store, etc.; • Measures revealing the load on each query dispatcher, which include, the number of queries that are currently active on the query dispatcher since the last index set change, the number of queries that have been dispatched since the last index set change, etc.; • Metrics that measure the load on each indexer partition, such as, the number of active items on every indexer partition, the number of items that are indexed per second, the current state of a partition, etc.; • Metrics revealing indexer health, such as, the total number of feed operations processed by each indexer, the total number of feed operations processed by an indexer, the current load on the API queue of the indexer, etc.; • Key measures of the processing ability of the Query & Results server, such as, revealing the rate at which the system queries and user queries failed, the number of queries handled by this QR Server per second, etc.; • Metrics revealing the root-cause of slowdowns in search queries, such as, the elapsed time between acknowledgement of a batch submission and the receipt of a final success or failure callback, the time taken from item initialization to completion of indexing, the time that elapsed between placing an item in a batch and its submission, the count of batches that are ready for submission, the ones that are submitted, items that are currently processed, etc.; • Measures revealing the current status of each document processor;

Microsoft Windows Applications	Windows Clusters	<ul style="list-style-type: none"> Metrics revealing the current status of the resource groups managed by the cluster service, such as, the number of groups that are offline/online/partially online; Key metrics that report the status of the nodes in the cluster, which include, the number of nodes that are active, the number that is down, and the number that is paused; The number of network interfaces in the cluster that are currently running and those that are currently down; The number of network interfaces in the cluster that are currently running and those that are currently down; The number of cluster resources that are online/offline, the number of cluster networks that are up/down, etc.;
Storage and Backup Applications/Devices	NetApp Filers	<ul style="list-style-type: none"> Key metrics pertaining to the NetApp filer hardware such as failed number of fans, failed power points, etc.; Statistics indicating the utilization of system resources such as the percentage of time for which the CPU was busy, the rate of data received and sent etc.; Performance statistics pertaining to the filer's disk drives such as the number of active, broken, and spare disks, etc.; Measurements pertaining to the filer's file system such as the disk space utilized and free, etc.; Statistics related to the RPC, NFS, and CIFS protocols such as the rate of calls received and rejected by the RPC, NFS, and CIFS layers respectively, etc.
	NetApp Netcache	<ul style="list-style-type: none"> HTTP, FTP, NNTP request related measures such as the rate of requests, rate of request hits, misses, response time for requests, etc. for each protocol; Measures related to streaming requests such as the rate of streaming requests, rate of request hits, misses, response time for requests, etc.
	NFS (Solaris only)	<ul style="list-style-type: none"> Server statistics related to RPC calls such as the total number of calls received, the number that was bad, etc.; Client Statistics related to RPC calls such as the total number of calls received, the number that was bad, etc.; The availability and access times of network file systems remotely mounted by a client
	Symantec Backup Server	<ul style="list-style-type: none"> Performance metrics pertaining to the server such as the number of active jobs, failed jobs, successful jobs and the size of data backed up; Key metrics such as the percentage of jobs that were aborted, corrupted, skipped, etc.; Metrics revealing the size of the VMware Virtual Machines that have been backed up, the number of virtual machines that have been backed up and the percentage of data backed up per virtual machine etc.,

	Veeam Backup Server	<ul style="list-style-type: none"> Metrics revealing the current status of the jobs executing on the Veeam backup server and the time taken for execution; Metrics revealing the current state of each job and the size of each job;
SAP Applications	SAP R/3 Servers	<ul style="list-style-type: none"> Key statistics relating to the SAP R/3 server's buffers such as the percentage of buffer space used, the buffer hit ratio, the swap rate, etc.; Measures indicating the R/3 server's memory management capabilities, which include the maximum stack memory usage, the percentage of stack memory actually used, the total heap size and percentage heap usage, the percent utilization of extended memory slots, etc.; Performance metrics pertaining to the various server components such as the spool service object, the background service, the dialog service, enqueue service, etc.; SAP database-related metrics, which include the extent of usage of the database tablespaces, the database buffers, etc.; SAP gateway-related information such as the percentage utilization of gateway connections, the percentage of remote gateways, etc.; Metrics reporting the occurrence of ABAP short dumps such as the total number of dumps and the newly introduced dumps; External measures of server availability and connect time; Key measures revealing the level of user activity on the SAP server
	SAP AGate Servers	<ul style="list-style-type: none"> Key component-level metrics such as the percentage utilization of user sessions and worker, the hit rate of the component, etc.; Important access-related statistics such as the number of times the AGate component was successfully accessed; Error statistics such as the number of errors that occurred in the AGate component during the last measurement period; Measures revealing the number of web accesses serviced by an AGate instance
	SAP BOBI	<ul style="list-style-type: none"> Metrics revealing the current health, status and thread usage of the Dashboard Cache Server; metrics revealing the request processing capability of the Dashboard Cache Server; Metrics revealing the current health, state and thread usage of the File Repository server; metrics revealing the request processing capability of the File Repository server; Statistics revealing the current health, state and thread usage of the Adaptive Job server; metrics revealing whether the Job server is able to send documents to various designated destinations; number of jobs received for processing by the job server and the number of jobs that failed; Statistics revealing the health, state and the processing ability of the Adaptive Processing Server; Statistics related to the JVM health and the correctness of the configuration of the critical Adaptive Processing Server services;

SAP Applications	<ul style="list-style-type: none"> Metrics revealing the status of the connection server; running state of the connection server; number of server threads that are currently servicing requests; Statistics revealing the current health and running status of the Crystal Reports Server; usage of the server threads; metrics revealing the rate at which the server processes the requests; Key metrics revealing the current health and running status of the Dashboards processing server; usage of the server threads; metrics revealing the request processing capability of the Dashboard processing server; Statistics revealing the current health, status and request processing ability of the Report Application server; Metrics revealing the current health, status of the Web Intelligence server; tracks the load on the server cache through its size; tracks client calls made to and sessions created on the Web Intelligence server; Logs related to the core server types such as the Central Management server, Adaptive Processing server, Adaptive Job server etc; Metrics revealing the requests received by and the connections to the semantic layer; availability and processing ability of the Central management server; metrics revealing whether the CMS has a healthy connection to the ADS; the count of the jobs that failed and are waiting in the CMS; metrics related to the connections established by the CMS and the utilization of these connections;; Metrics revealing the number and type of sessions that are currently active on the CMS; license utilization of the CMS etc; Statistics revealing the availability and thread pool size of the Event server; Status of the Web Application Container server; tracks Multi-Dimensional Analysis Service sessions and reports the number of OLAP data requests received; Metrics related to the availability of the Platform Search service; status of the indexing mechanism; document indexing rate etc; Statistics revealing the availability of the Data Federation Service; the load in the service, the query load on the data federation query engine etc; Metrics revealing the current health of the monitored node, execution time of each probe, current state of each service running on the node, utilization of each service operating in the monitored node etc;
SAP Web Application Server	<ul style="list-style-type: none"> Measures relating to the Configuration Manager's interactions with the database, such as the rate of database lock exceptions, the percentage of cache reads, etc.; Key statistics pertaining to application and system threads, which include thread pool usage, the number of tasks waiting to be executed, etc.; Performance metrics revealing how well client connections are managed, which includes the current connection pool size, the number of unrecognizable connections, etc.; Pool manager-related measures, such as the memory allocated to the pool manager, and percentage of memory utilized; Measures revealing the availability and responsiveness of the P4 connection to the server;

		<ul style="list-style-type: none"> Key statistics pertaining to how well the EJB container service manages the enterprise bean instances, which include the number of bean creations/removals, the number of bean passivations and activations, etc.; Session-oriented performance measures, which include the number of failed logon attempts, the number of invalid, logged off, and timed out sessions, etc.; Transaction-related statistics, such as the number of rolled back transactions, suspended transactions, etc.
SAP Applications	SAP Web Dispatcher	<ul style="list-style-type: none"> Key statistics indicating the current state, number of connections and maximum time taken to establish the connection. Measures revealing the rate at which the requests are processed by the thread and status of the thread. Threads related measures such as total number of threads created from the thread pool, ratio of the total number of threads created to the maximum thread setting, number of currently open connections, etc. Metrics revealing current load on the destination, current number of HTTP and HTTPS connections, stateless and stateful requests to the destination, and average time taken by a request and recent response time for a ping request to the destination. Measures revealing number and percentage of expired or invalidated entries in the cache.
	Max DB	<ul style="list-style-type: none"> Basic measures of availability and responsiveness of the database; Memory usage statistics such as the percentage of memory in data area that is utilized; Locking-related metrics, which include, the rate of deadlocks and collisions, the number of requests awaiting locks, etc.; Usage metrics pertaining to the log queue, such as, the maximum number of transactions written to the queue, the log queue overflows, etc.; Performance metrics pertaining to the SAP liveCache, which include, the memory in use, the number of attempts to acquire a spinlock, the rate of errors, etc.; The hit rate of the data cache and the database session cache; I/O buffer cache-specific metrics, such as, the current size of the cache, the percentage of the cache used by the data cache and the converter, the free space in the cache, etc. Transaction-related statistics, which include, the rate of transaction commits and rollbacks, the rate at which SQL commands were executed and parsed, etc. Metrics revealing the type of queries that are executed on the database, and the number of active sessions to the database

Corillian's Voyager	Voyager Transaction Processor (TP)	<ul style="list-style-type: none"> Statistics pertaining to the authentication database and the TP database such as the database pool size, the rate at which requests were serviced/enqueued, the average transaction response time etc.; Measures that monitor the requests to the TP such as the total number of requests processed by the TP, the average response time of the TP, etc.; Statistics revealing the health of the authentication service such as the rate of database calls, the number of authentication API calls, database API calls, identity API calls, pin vault API calls that are currently in progress, etc.; Metrics relating to the TP sessions such as the total number of Voyager sessions on the TP, the sessions that have timed out, etc.; Performance statistics pertaining to the host server such as the number of transactions inside a host server, the transaction processing time of the host server, etc.; VLB-related metrics which include the number of requests executing on the VLB, the number of requests processed by the VLB, the request execution time, etc.
	Voyager User Interface	<ul style="list-style-type: none"> Measures such as the number of executions, the transaction processing time, the rate of incoming/outgoing data, etc.
Virtualization Software	VMware® vSphere/ESX Server	<ul style="list-style-type: none"> A single eG agent performs a patent-pending 'In-N-Out' monitoring to extract the percentage of physical CPU, memory, disk, storage, and network resources used by the ESX host, the percentage of physical resources used by each of the guest operating systems executing on the host, and the percentage of 'allocated' resources used up by each of the guests. Reports on status of every guest in terms of the number of guests powered on and off, added and removed guests, etc.; for 'virtual desktops', the number of guests with and without users is also reported; Metrics specific to 'virtual desktops', such as, the number of new logins to the desktop, the number of sessions logging out, etc.
	Microsoft Virtual Server	<ul style="list-style-type: none"> A single eG agent on the base operating system reports the percentage of physical resources (CPU, memory, disk space, etc.) consumed by the virtual host, and polls each of the guests to determine the fraction of physical resources used up by every guest. Monitors the event logs on each guest to trap application/system errors and warnings Reports metrics indicating the status of guests – i.e., whether powered on or not; in-out movement of guests can also be tracked using metrics such as the number of guests that were added/removed during the last measure period

Virtualization Software	Solaris Zones	<ul style="list-style-type: none"> • A single eG agent on the base Solaris host reports the percentage of physical resources (CPU, memory, disk space, etc.) consumed by the virtual host, and polls each of the zones to determine the fraction of physical resources used up by every zone. • Reports metrics indicating the status of zones – i.e., how many are registered, installed, running, etc.; in-out movement of zones can also be tracked using metrics such as the number of zones that were added/removed during the last measure period • Monitors the uptime of every zone configured on a Solaris virtual host
	Solaris LDom s	<ul style="list-style-type: none"> • A single eG agent on the primary domain reports the percentage of physical CPU/memory resources allocated to the primary domain and other logical domains, the percentage of physical CPU/memory resources utilized by the primary domain and the Solaris host as a whole, and also the percentage of allocated resources consumed by the primary domain. • Reports metrics indicating the status of the logical domains – i.e., how many are registered, active, inactive, in bind/unbind state, etc.; in-out movement of domains can also be tracked using metrics such as the number of domains that were added/removed during the last measure period • Provides an “outside view” that reveals the percentage of physical CPU/memory resources utilized by each logical domain; • Provides an “inside view” that reveals how well each domain consumes the allocated CPU/memory/disk resources, the network and TCP traffic to and from every logical domain, and the uptime statistics per domain;
	Citrix XenServers	<ul style="list-style-type: none"> • A single eG agent on the control domain reports the percentage of physical CPU resources used up by the control domain and each of the processors that the XenServer host supports, the total disk capacity and usage of each volume group of the host, the physical memory consumed by the control domain, the network traffic to and from the host, etc.; • Reports metrics indicating the status of the guests – i.e., how many are registered, running, halted, suspended, etc.; in-out movement of guests can also be tracked using metrics such as the number of guests that were added/removed by XenMotion; • Provides an “outside view” that reveals the percentage of physical CPU/memory resources utilized by each guest; • Provides an “inside view” that reveals how well each guest utilizes the allocated CPU/memory/disk resources, the network and TCP traffic to and from every guest, and the uptime statistics per guest;
	Citrix Provisioning Servers	<ul style="list-style-type: none"> • Measures revealing the availability of the License server and the database server used by the Citrix PVS; • Key metrics revealing the composition of a PVS farm, such as the number of sites, servers, stores, and farm views in the farm; • Status of the Citrix PVS;

Virtualization Software	Citrix Provisioning Servers	<ul style="list-style-type: none"> • Site-specific statistics reporting the number of active/inactive servers, devices and vDisks in each site; • vDisk-related metrics, which include, the status and size of each vDisk, whether the vDisk is currently locked or not, etc.; • Statistics pertaining to device collections, such as, the number of active/inactive devices in each collection, etc.
	VDI in a Box / VMware / XenServer / Hyper-V	<ul style="list-style-type: none"> • License related statistics such as the licenses of each type that are currently installed, number of licenses currently in use, the available licenses and the usage statistics of each license type etc., • Metrics revealing the current status and type of the server, RAM size, CPU cores etc., desktop related statistics such as the number of desktops used, desktops that are prestarted etc., space usage statistics such as the total space, free space available in the logical storage of the server etc., • Key metrics revealing the number of processes executing on the server, their CPU and memory utilization; Availability and responsiveness of the configured TCP ports on the server; • Template related statistics such as the desktops created from each template, desktops that are currently in use, desktops that are currently in broken state;etc., • Status of the user's desktop sessions • Status of the image, image distribution across the servers of the grid, number of templates that are currently using the image, is the HDX protocol enabled?, is the desktop agent installed etc., can be determined • Key metrics revealing the status of the template, refresh policy configured for the desktops, RAM allocation and virtual cores allocation to the desktops, maximum number of desktops that can be generated from the template and the desktops that are already started and are ready for login etc.,
	KVM Servers	<ul style="list-style-type: none"> • Measures relating to the virtual components such as the sockets, threads, cores, virtual CPUs etc, • Measures relating to the physical memory allocation of the KVM server host; • Measures revealing the status of each storage pool in the KVM server and the space utilization of each storage pool; • Measures revealing the type of each storage volume and the space utilization of each storage volume; • The current status of the virtual network; • Measures relating to the resource utilization of each of the virtual machine hosted on the server, overall status of the virtual machines etc; • Measures relating to the data/file processing during live migration of the VM; • Apart from the above measures, the eG agent collects a host of metrics pertaining to individual virtual guests such as the resource utilization of each virtual machine, TCP traffic, network loading etc,



Measurements made by eG Agents

Virtualization Software	Citrix XenMobile MDM	<ul style="list-style-type: none"> The current status of the Citrix XenMobile MDM License utilization details of the Citrix XenMobile MDM such as the total licenses used, licenses held by the XenMobile MDM, number of days by which the license will expire etc; Key statistics indicating the status of the scheduled jobs such as the jobs that are currently running, cancelled jobs, recent jobs that failed, jobs that ran recently etc; Thread related statistics such as the maximum number of threads the server can spawn, the threads that are currently in use, the waiting threads, the queue length, threads that were cancelled etc; Key statistics mentioning the number and names of the devices that are using hosting blacklisted applications, devices that are hosting applications that are not suggested, devices that are missing the installation of suggested applications etc; Numerical statistics revealing the full wipes, pending full wipes, corporate wipes completed, device locks completed, tracks completed etc; The current load on the server, the count of the devices currently connected to the server, resume requested devices etc; Operating system related statistics such as the devices that are currently managed/unmanaged, devices that are currently active/inactive, devices that have violated one/more policies etc; Key statistics revealing the status of the devices connecting to the corporate network such as the currently managed/unmanaged devices, active/inactive devices etc; Package related statistics such as the pending deployments, successful deployments, failed deployments etc;
	Citrix XenMobile	<ul style="list-style-type: none"> Memory details of the Citrix XenMobile such as the memory allocated to each memory pool at startup stage, upper limit usage of initially allocated memory, amount of memory that is currently utilized by each memory pool, amount of committed memory, maximum amount of memory used for memory management by each memory pool, etc.; JVM threads related statistics such as total number of threads (including daemon and non-daemon threads), highest number of live threads since XenMobile JVM started, number of threads in different states such as runnable, blocked and timed waiting states, etc; Metrics revealing number of times the collections were loaded to, fetched from, recreated in the Hibernate Cache etc; Measures indicating number of times the entities were loaded to, fetched from, inserted into the Hibernate Cache etc; Statistics shed light on the number of queries that were successfully retrieved from the queries cache, average time taken to execute the queries in the queries cache, number of queries executed from the queries cache etc;

Virtualization Software	Citrix XenMobile	<ul style="list-style-type: none"> • The number of cacheable entities/collections that were successfully retrieved from the second level cache, number of cacheable entities/collections stored in the memory of the second level cache, total size of cacheable entities etc.; • Key measures mentioning total number of JDBC connections requested by the sessions, number of sessions that were opened/closed, the time taken to execute the slowest recorded query etc; • Key statistics mentioning maximum amount of memory allocated for JVM, amount of memory that is currently in use, time taken by the garbage collector for collecting unused memory etc; • SSL certificate related details such as the current status (whether valid/invalid) and how long this certificate will remain valid. • Measures revealing current status of each cluster task performed in the XenMobile and time duration elapsed since the cluster task was last updated. • Current status of each cluster node in the XenMobile cluster and the total number of connections made to each cluster node. • Statistics indicating number of successful, pending and failed deployments of each device policy and each delivery group in the XenMobile environment. • The current workload imposed by actions on the XenMobile. • Number of devices that are currently in session state. • Measures shed light on the number of messages received for each action, number of messages received owing to cancellation of this action, etc; • Metrics indicating total number of users currently connected to the XenMobile server, percentage of users who logged in recently etc; • License usage metrics such as the number of licenses of each product category that are currently in use, the number and percentage of licenses utilized, etc. • Measures indicating connection status of the user with the Xenmobile, time taken to connect to XenMobile and authenticate user login, etc; • Measure revealing the total number of connections made to each cluster node in the XenMobile. • XenMobile threads related statistics such as the maximum number of threads the server can spawn, the threads that are currently in use, the waiting threads, the queue length, threads that were cancelled etc; • Statistics revealing that how many times each operation succeeded and failed. • Key statistics indicating the number of jobs that keep repeating, total number of jobs that were scheduled, the jobs that are currently running, cancelled jobs, recent jobs that failed, jobs that ran recently etc; • Current status of each connection established to XenMobile. • Key measures mentioning the number of jail broken devices, number of devices that are newly enrolled, number of non-compliant devices, number of devices of this type that are owned by the company and the employees, number of devices that are managed/unmanaged by the server, etc;

Virtualization Software	Citrix Sharefile	<ul style="list-style-type: none"> The availability of the control plane, time taken to connect to the control plane, authentication status, time taken for authentication etc; The availability and access time of each storage zone configured on ShareFile, the number, names and types of the storage zones; License utilization details of the ShareFile such as the number of employee licenses, used licenses, usage etc; Key metrics revealing the storage space utilization of the ShareFile; Statistics revealing the devices that are currently active on the ShareFile account, load on the ShareFile, wiped devices, locked devices etc; Metrics revealing the number of files and folders in the ShareFile account, size of the folder, active folders, least used folders etc; Key metrics revealing the status of uploads, downloads and deletes performed using ShareFile, time taken to upload/download/delete the file etc;
	Citrix AppController	<ul style="list-style-type: none"> Key metrics that capture the expiry date of all the active SSL certificates, metrics that help in computing the validity of the active certificates etc; Measures revealing whether an user can connect to the AppController, time taken to connect, status of user authentication during login, time taken to authenticate user logins, time taken to login etc; Key metrics revealing the user operations that succeeded and failed; Numerical statistics revealing the key usage policies enforced on applications of each type; numerical statistics revealing the successful and failed application launches; Metrics revealing the number of user sessions in the AppController, the number of users who logged in recently, the sessions that logged out and the sessions that failed; Key metrics revealing the number of sessions that are open for each user, the number of successful and failed application launches for each user; Numerical statistics revealing the devices that are connected to the AppController, devices that are locked and devices that are erased; Measures revealing the number of users currently logged in through each received connecting to the AppController, number of users from the internal network who logged into the AppController through this receiver, number of users who logged into the AppController through an external network.
		Microsoft Hyper-V

Virtualization Software	<ul style="list-style-type: none"> • Key metrics revealing how well the root partition manages the host's physical resources, which includes, the number of address spaces in the TLB of the partition, the number of virtual processors in the root partition, the number of pages currently deposited into the root partition, the number of page tables that are currently present in the virtual TLB of the root partition, the rate of flushes of the entire TLB etc. • Metrics revealing virtual processor usage of the parent partition, such as, the percentage of time the virtual processor of the parent partition spent in executing guest code and in executing hypervisor operations, the percentage of time the virtual processor was in use, the rate at which CPUID, emulated, HLT instructions were completed, the rate of small and large TLB fills, etc. • Metrics measuring the health of the hypervisor, which includes, the number of virtual processors supported by the hypervisor, the number of partitions managed by the hypervisor, the number of bootstrap and deposited pages in the hypervisor, etc. • Statistics indicating how Hyper-V utilizes its network adapters, switches, and ports, such as, the rate of broadcast packets sent/received, rate of data sent/received, etc. • Reports metrics indicating the status of the guests – i.e., which VMs are powered-on/off; in-out movement of guests can also be tracked using metrics such as the number of guests that were added/removed; • Provides an “outside view” that reveals the percentage of physical CPU/memory/disk resources utilized by each guest; • Provides an “inside view” that reveals how well each guest utilizes the allocated CPU/memory/disk resources, the network and TCP traffic to and from every guest, and the uptime statistics per guest; • A specialized 'Hyper-V VDI' model that reveals the users logged into the virtual desktops on the Hyper-V server, and the resource usage of each user; the total number of currently open sessions across the virtual desktops, the percentage of new logins to the desktops, the number of sessions that logged out, etc.
AIX LPARs on IBM pSeries Server	<ul style="list-style-type: none"> • Metrics revealing the overall physical CPU usage by all AIX LPARs, such as, the number of installed processing units on the IBM pSeries server, the number of free processing units on the server, the number and percentage of used processing units, the percentage of physical processors consumed by AIX LPARs, the number of dedicated and shared processors, etc. • Metrics measuring the I/O activity on each physical adapter on the Virtual I/O server, such as, number of data transfers handled by each adapter, the amount of data received and sent by each adapter, etc. • Metrics indicating the composition of volume groups on a Virtual I/O server and how the volume groups are utilized, such as, the status of each volume group, the number of logical and physical volumes in each volume group, the number of volumes that are currently active in each group, the number of allocated and unallocated partitions in each group, etc.

		<ul style="list-style-type: none"> Metrics that alert you to space crunches in storage pools, which include, total storage pool size, allocated pool size, percent usage of space in storage pool, percent free space in storage pool, etc. Reports metrics indicating the status of the AIX LPARs – i.e., which LPARs are powered-on/off; in-out movement of AIX LPARs can also be tracked using metrics such as the number of LPARs that were added/removed; Provides an “outside view” that reveals the percentage of physical CPU/memory utilized by each AIX LPAR; Provides an “inside view” that reveals how well each AIX LPAR utilizes the allocated CPU/memory resources, the network and TCP traffic to and from every AIX LPAR, and the uptime statistics per LPAR;
Virtualization Software	RHEV Hypervisor	<ul style="list-style-type: none"> Measures indicating the hypervisor status, physical CPUs available to the hypervisor, etc. Metrics revealing the memory usage of the hypervisor host, such as, the used physical memory, percentage of free physical memory, swap memory usage, memory overcommitment, etc. Physical CPU usage-related metrics, such as, the percentage of user CPU, system CPU, and idle CPU utilization; Metrics revealing network health, which include, the current status of each network interface, the traffic on and errors experienced by each network interface, etc. Metrics revealing how efficient the virtual network is, which include, the rate of data transmitted and received over the network, and the errors encountered in the process; Outside-view metrics revealing how each VM uses the physical resources of the host, such as, power-on state of each VM, the virtual CPUs allocated to and utilized by VMs, physical memory consumption, disk throughput, network bandwidth consumed by VMs, etc. VM status-related metrics such as the number of registered, powered-on, powered-off, suspended, orphaned VMs, etc. Inside-view metrics revealing how each VM uses the allocated resources;
	RHEV Manager	<ul style="list-style-type: none"> Errors/warning events captured by the RHEV manager log; Storage domain-related metrics, such as, the number of VMs using a storage domain, the capacity of and space used on a storage domain, storage domain availability, etc. Metrics revealing the number of information, error, and warning events that occurred on the manager; Logical network-related metrics, such as, the current status, bandwidth used, and errors experienced by logical networks;

		<ul style="list-style-type: none"> Cluster-related metrics, which include, the Physical CPU available to cluster, the physical CPU usage of VMs in cluster, physical memory usage of VMs in cluster, total physical hosts in cluster, powered-on/power-off hosts in cluster, VMs in cluster, etc. Datacenter-related metrics, such as, the current status of a datacenter, the disk capacity, free space in datacenter, clusters, servers, VMs in datacenter, etc.
Virtualization Software	HMC Server	<ul style="list-style-type: none"> Metrics revealing the availability of HMC for remote command execution and response time which indicates the duration for which the HMC was connected; Memory related metrics such as the total memory, used memory, free memory, swap memory and the percent of used memory; Metrics relating to Resource Monitoring and Control (RMC) tasks and HMC tasks; each indicating the total number of tasks, number of running tasks, number of sleeping tasks, number of stopped tasks and the number of zombie tasks; Processor metrics that include CPU utilization, percent of used system CPU, percent of idle CPU, percent of IO Waits and the total number of zombie HMC tasks; Reports metrics that indicate the number of pSeries servers that are managed in the HMC server and the number of LPARs that are available in the managed pSeries server; Metrics relating to the Web logins and the terminal logins; each indicating the number of current sessions, new sessions, percent of new sessions and the number of disconnected sessions from the server; Metrics relating to the uptime of the servers such as rebooted, measure indicating the time period during which the system has been up since the last measurement and the total time during which the server has been up since the last reboot.
	VMware vCenter	<ul style="list-style-type: none"> Metrics revealing the availability and responsiveness of vCenter; Session statistics that indicate the current session load on VirtualCenter, sessions that logged out suddenly, new logins to VirtualCenter, etc.; Key metrics that measure how effectively VirtualCenter manages ESX server licenses, which include, the number of licenses installed, the number and percentage of licenses utilized, etc.; Statistics that indicate the current configuration of ESX server clusters managed by VirtualCenter, and the CPU/memory resources used by each cluster; The resources available to every resource pool under each ESX cluster managed by VirtualCenter, and the extent to which the resource pool utilized the available resources;

Virtualization Software	Oracle VirtualBox	<ul style="list-style-type: none"> Metrics revealing host health, such as, the total physical memory, physical memory usage, free physical memory on the host, CPU usage of the host, percent CPU usage on user, system-level, and idle processing, etc. Statistics revealing the health of VMs operating on VirtualBox, such as, the powered-on state of a VM, the RAM, video memory, and balloon memory configuration of a desktop, RAM usage, CPU usage, etc.; Session-related metrics, such as, current sessions, percentage of new logs, sessions logging out, etc.; Metrics revealing VM status, such as, the number of registered, running, not running, added, removed guests; Inside-view metrics of every desktop, which include, the user who is currently connected to a desktop, the time and duration of connection, and how the user utilized the resources allocated to every
	Oracle VDI Broker	<ul style="list-style-type: none"> Status of the RDP Broker service; Number and nature of errors/warnings logged in the Oracle broker's log file; Metrics revealing the administrative, operational, and availability state of the broker; Current state and resource usage of the Oracle VDI manager; Desktop provider-related metrics, which include, the current state of the desktop provider, number of pools and datacenters managed by the desktop provider, the percentage of provider's desktops assigned to users, the percent CPU, memory, and storage usage of provider's desktops,, etc.; Metrics revealing the administrative, operational, and availability state of the broker; Pool-related measures, such as, the assignment status of pools, the type of desktops assigned to each pool, total desktops in pool, the number of running, powered off, unknown, suspended desktops in pool, etc.; Desktop-related metrics, which include, the number of desktops managed by the broker, the number of running, powered off, unknown, suspended desktops, etc.; Per desktop metrics, which indicate, the machine and desktop state of a desktop, the amount of RAM, disk capacity, and video memory allocated to a desktop, etc.
	Oracle VM Manager	<ul style="list-style-type: none"> Key metrics revealing the count and details of the critical information, error and warning events that are generated on the Oracle VM Manager; Numerical statistics that reveal the jobs that started, completed, failed, succeeded etc; the average time taken to complete the jobs; the number of jobs that were outstanding in the job queue etc; The current state of the Oracle VM Manager; the availability of the web interface and the time taken to connect to the web interface;

		<ul style="list-style-type: none"> Key metrics revealing the number of registered Oracle VM servers in each pool, the number of servers that are currently running, number of servers that are not running and those that are in maintenance mode, the number of VMs in each pool and the VMs that are currently running etc; The total capacity and space utilization related metrics for each file system of the Oracle VM Manager; metrics revealing whether the file system was shared or not; The total capacity and space utilization related metrics for each SAN storage used by the Oracle VM Manager etc;
Virtualization Software	Oracle VM Server	<ul style="list-style-type: none"> Key metrics revealing the physical CPU usage of the Oracle VM server such as the number of processors, number of cores per socket etc; the speed at which the processors have been configured to run; Numerical statistics revealing the processors assigned to the control domain; the CPU utilization and memory utilization of the control domain; Key metrics revealing the physical memory utilization of the VMs and the total amount of physical memory on the host; The total capacity and space utilization related metrics for each file system on the Oracle VM Server; metrics revealing whether the file system was shared or not; The total capacity and space utilization related metrics for each SAN storage used by the Oracle VM Server etc; Key metrics revealing the CPU time utilization of the control domain, run queue length, swap memory of the control domain etc; Metrics revealing the space utilization of each disk partition of the control domain; the input/output utilization of each physical disk etc; The current status of the Oracle VM Server, metrics revealing whether the server is in maintenance mode, the availability of the web interface and the time taken to connect to the web interface; The current state of each VM on the Oracle VM Server, numerical statistics of the physical server's resources that each VM on the Oracle VM server is taking up; Numerical statistics revealing the VMs that are registered, powered on, added and removed; Provides an "inside view" that reveals how well each VM utilizes the allocated CPU/memory/disk resources, the network and TCP traffic to and from every VM, and the uptime statistics per VM etc;
	VMware vCenter Cluster	<ul style="list-style-type: none"> Metrics revealing the availability and responsiveness of vCenter cluster; Session statistics that indicate the current session load on vCenter cluster, sessions that logged out suddenly, new logins to vCenter cluster, etc.; Key metrics that measure how effectively the cluster managers ESX server licenses, which include, the number of licenses installed, the number and percentage of licenses utilized, etc.; Statistics that indicate the current configuration of ESX server clusters managed by the vCenter cluster, and the CPU/memory resources used by each ESX cluster;

		<ul style="list-style-type: none"> The resources available to every resource pool under each ESX cluster managed by the vCenter cluster, and the extent to which the resource pool utilized the available resources;
Virtualization Software	Docker	<ul style="list-style-type: none"> Metrics indicating whether/not the Docker service is installed, loaded or running currently. Measures revealing percentage of data space and Metadata space utilization; Key metrics disclosing the numerical statistics of various events that occur in the Docker. Docker image related metrics such as total number of images, number of images that are mapped/unmapped to the containers and disk space utilized by the mapped/unmapped images. Measures indicating the disk space utilization of each image, number of containers created from each image, number of containers (created from each image) that are running currently, amount and percentage of memory utilized by the containers that were created by each image etc; Metrics that shed light on the Docker containers such as total number of containers, the current status (whether/not running) of the containers, number of containers that are newly added to or removed from the Docker server, memory utilization of the containers etc.; Key measures revealing the uptime of each container, packets transmission, the rate at which data is transmitted from each container, incoming and outgoing traffic, CPU utilization of each container, etc.; Statistics indicating average network delay during packet transmission, packet loss and availability network connection;
	Nutanix Acropolis	<ul style="list-style-type: none"> Key measures indicating the number of times the read requests served from each cache, the real memory consumed by the data in the cache, the amount of memory saved due to deduplication, the logical SSD memory used to cache data without deduplication, etc; Measures revealing the current status, type and mode of each physical disk, whether the data in the disk is migrated or not, the total capacity of the disk, the amount of the disk that is currently in use and is still unused, the total number of I/O operations that are currently performed on the disk, the time taken for processing the I/O requests, and bandwidth used by the disk when processing I/O requests; Metrics reveal the average time taken by the storage to process read and write I/O requests, the bandwidth utilized for processing the read and write I/O requests, the number of read and write operations that are currently performed on the storage, the total capacity of the storage, the amount of storage space that is still unused in the SSDs and directly attached SATA HDDs, etc;

Virtualization Software		<ul style="list-style-type: none"> System related measures such as number of CPU cores and CPU sockets on the host, the total capacity of the host across all the CPU cores, total memory capacity, the current size of the oplog, etc; Measures revealing the logon and logoff duration for each user, size of each user's profile when it is retrieved from the user's store at logon, the number of locally copied files that are synchronized during logon and log off and categorized by the file size of 1KB and the file size ranging from 1KB to 10KB, 100KB to 10MB, 1MB to 5MB and 5MB and above, etc; Metrics indicating the current state of the virtual machine, whether the VM is the controller VM or not, the number of sessions that are currently active on the virtual machine, percentage of virtual and physical CPU resources used by the VM, the amount of data received and transmitted by the VM, etc; Statistics revealing the number of VMs that are currently powered on, the number of powered off VMs on the hypervisor, the number of suspended VMs, the number of powered on guests that are currently logged in with users, the number of that are newly added to and newly removed from the Nutanix Acropolis; Measures reporting the number of sessions that are currently active across all the guests, the number of new logins, the percentage of current sessions that logged in and the number of sessions that logged out.
	Nutanix Prism	<ul style="list-style-type: none"> Measures indicating the replication factor of each container, the replication factor setting for the Oplog, whether on-disk deduplication and erasure coding are enabled or not, the number of VMs that are attached to the container, maximum capacity configured for the container, etc; Metrics revealing the availability of the prism, the time taken by the Prism service to respond to HTTP requests and the response code returned by the emulated HTTP request; Storage pool related metrics such as the number of disks pooled in the storage pool, the amount of space in the cluster that is available to the storage pool, the total amount of logical storage space used, the amount of actual usage of storage (i.e., usage after compression and deduplication), etc; Measures indicating whether the cluster is deployed on the cloud or not, whether shadow clones are enabled for the cluster or not, whether/not the lockdown mode has been enabled for the cluster, the number of hypervisors, the average time taken by the physical disks to process read and write I/O requests, the total storage capacity, etc; Statistics revealing the total number of VMs in each cluster, the number of VMs in the cluster that are currently powered on and are currently powered off, and number of controller VMs in the cluster; Metrics reporting the amount of data read from and written to the cluster in response to the read and write I/O requests, total memory capacity of the cluster, the amount of memory that is utilized and is unused in the cluster;

Connection Brokers	VMWare Virtual Desktop Manager (VDM)	<ul style="list-style-type: none"> Key metrics revealing the health of the network connection to the VirtualCenter, and the availability of the VirtualCenter; Statistics revealing if any error entries have been logged in the ADAM event logs; Metrics indicating the health of the Tomcat JVM, such as, the memory usage of the JVM, the free memory on the JVM, the total number of daemon and live threads on the JVM, the JVM uptime, etc.
	Citrix XenDesktop Broker 3/4	<ul style="list-style-type: none"> Metrics revealing the health of the IMA- interactions between a DDC and other servers in the farm, which include, the number of active IMA connections, the rate at which data was received and sent by DDC; Metrics revealing the health of the Citrix XenApp technology on which the DDC has been built, which include, the rate of application enumerations; the rate of application resolutions, whether a connection to the datastore is available or not, the rate of data written to and read from datastores and dynamic stores. Metrics indicating the availability of the servers hosting virtual desktops in desktop groups; Metrics indicating the errors recently captured by the DDC logs; Metrics revealing the availability and responsiveness of the DDC; Metrics revealing the status of virtual desktops in farms and desktop groups, such as, the number of desktops that are powered on/off, the number of desktops that have been assigned to users, the number of desktops in use, the number of desktops that are idle, etc.; Metrics reporting the composition of a DDC farm, which include, the number of DDCs in the farm, the number of desktop groups in the farm, the number of desktop groups that are available/unavailable; Metrics revealing whether the virtual desktop agent is available or not on each of the virtual desktops managed by DDC; Metrics indicating whether the virtual desktops are available over the network or not;
	Citrix XenDesktop Broker 5	<ul style="list-style-type: none"> Metrics revealing the connection status between the XenDesktop broker and each server hosting the virtual desktops; Metrics reporting the current status of each controller and the number of desktops registered with each of them; Metrics indicating the health of the interactions between the Citrix Broker Service, AD Identity service, Configuration server, Machine Creation service, Machine Identity service, and the database, such as, the database availability, the time taken for database transactions, transaction errors, etc.; Key metrics related to the desktops available per catalog, such as, the allocation type of desktops, the number of desktops assigned/not assigned to users, the number of desktops not in groups, type of catalog, etc. Metrics revealing the availability of the License server, database server, etc.;

		<ul style="list-style-type: none"> Desktop group-related metrics, which include, whether the desktop group is available/not, the total number of desktops in each group, the available desktops, disconnected desktops, etc.; Metrics tracking the disconnected sessions to virtual desktops, such as, the total number of disconnected sessions, new disconnects, etc.; Metrics tracking the number of sessions accessing applications published on each virtual desktop; Whether the virtual desktop agent is available/not per desktop; Metrics providing status updates on virtual desktops, such as, the current desktop state, whether desktop is available over the network or not, etc.
Connection Brokers	XenDesktop Broker (v7)	<ul style="list-style-type: none"> Metrics revealing the status of the connection between the broker and the hosting server, whether the connection between the broker and hosting server is in maintenance mode, whether machine creation service is used to create provisioned machines etc., Site level details such as the availability of the license server, response time of the license server, active sessions on the site, brokers in the site, etc., Broker level details such as the state of the broker, machines registered with the broker, status of the services such as the broker service, AD Identity service etc., on each broker, connectivity between each service and the database, health of the transaction performed by each service on the database etc., For both Site and individual brokers, delivery groups related metrics such as the availability of the delivery group, total number of machines in each group, the available machines, used machines, disconnected machines, machines on which error was detected etc., For both Site and individual brokers, catalog level metrics such as the number of machines, user assigned machines, registered machines, unregistered machines, total user sessions, established sessions etc., allocation type of the machines, machines available to users from each catalog to delivery groups, total machines in catalog etc., Application level statistics such as whether the application is enabled or not, visibility of the application, number of currently running instances of each application, priority level set for resource usage of each application etc., For both Site and individual brokers, session related metrics such as the total sessions, active sessions, connected sessions, sessions in Unknown state etc., For both Site and individual brokers, metrics revealing the total logins and logouts for each protocol type, currently active sessions etc.,
	Citrix XenDesktop Director	<ul style="list-style-type: none"> Metrics revealing the status of the desktop OS machines in each delivery group configured in a site; Statistics relating to the number of machines of each type that are currently in a state of failure; current state of each Server OS machine, resource utilization of each machine and the session load on each machine; Metrics revealing the user connections to each delivery group, time taken for users to access the desktops/applications delivered by each group etc;

		<ul style="list-style-type: none"> Metrics revealing the user session related statistics of each delivery group in a site, the number of user connections to each delivery group in a site, the connection failures etc; Numerical statistics revealing the users who are connected to the machines/applications in each delivery group in a site, login duration of each user to access the desktops/applications, number of times each user has logged on etc;
Connection Brokers	Citrix XA/XD Site 7.x	<ul style="list-style-type: none"> Key metrics revealing the status of the connection between the XenDesktop broker and each server hosting the machine, Metrics revealing the non availability of the license server on the site, responsiveness of the license server inactive brokers and the session overload on the site; Key metrics revealing the current status of each delivery controller configured within a site and the number of machines registered with each controller; Statistics revealing the health of each controller service on the target controller, status of desktop OS machines in each delivery group configured in a site, statistics relating to the number of machines of each type that are currently in a state of failure; Metrics revealing the session load, resource utilization and current state of each Server OS machine; Catalog related statistics such as the catalog type, the type of desktops allocated to each catalog, the availability, usage and assignment of desktops within a catalog etc; Numerical statistics revealing the users who are connected to the machines/applications in each delivery group in a site, login duration of each user to access the desktops/applications, number of times each user has logged on etc; Metrics revealing the user session related statistics of each delivery group in a site, the number of user connections to each delivery group in a site, the connection failures etc;
	Leostream connection broker	<ul style="list-style-type: none"> External metrics revealing the availability and responsiveness of the connection broker and its web server component; Metrics indicating whether the internal and external databases of the connection broker are available or not, and if so, how quickly they respond to requests; Measures revealing the current status of the centers from which the connection brokers gather environment information; Job queue-related measures such as, the number of waiting jobs in queue, the number of finished, aborted, cancelled jobs, additions to the queue in the current measurement period in terms of new waiting, finished, in progress, aborted, and cancelled jobs; Statistics reporting whether the SunRay servers with which the connection broker is interacting is alive or not; Metrics indicating the number of registered desktops on the connection broker, the number assigned to users, the number of desktops assigned but unused, the number of free desktops, etc.

Connection Brokers	VMware View	<ul style="list-style-type: none"> Metrics revealing the status of the connection broker; Metrics that track license usage, such as, the number of days to the expiry whether/not the local mode license is enabled, etc.; Connectivity-related metrics that report whether/not the View Manager is able to connect to vCenter, whether/not the View server is able to connect to the Events Database, whether/not each virtual desktop (in a group) is available over the network, etc.; Reports connectivity between the vCenter and the View manager; also indicates the connectivity of this server with the events database, Active Directory; Usage metrics per desktop pool, such as, the total number of desktops, the number of active/inactive desktops, percentage of inactive desktops, etc. Metrics revealing the number of broker events, admin events, transfer events and agent events recorded in the events database;
	VMware Horizon View RDS	<ul style="list-style-type: none"> Metrics revealing the availability of the VMware view RDS connection and the time taken by the server to respond to a request; metrics revealing whether the user login process was successful and the time taken for a login; Metrics revealing the activity of redirector component of the Microsoft windows operating system and metrics revealing the status of the file serving as seen by a file server's client; the status of logons from client sessions to the VMware RDS server; Key measures revealing the time taken for user login and profile loading process; the number of times the user profile was loaded; the number of times the profile loading for each user was successful and the number of times the profile loading failed etc; Metrics revealing the number of instances of the published application executing on the server, the memory and CPU utilization of each published application; The status of the listener port, performance statistics relating to the VMware RDS server user sessions, metrics revealing the resource utilization of each user on the RDS server etc; Statistics revealing the number of disconnected VMware RDS server sessions, the new logins to the VMware RDS server etc;
Hardware Monitoring	Sun, HP, Dell, IBM hardware monitoring	<ul style="list-style-type: none"> Integrates with Dell Open Manage and HP/Compaq Insight Agents Measures hardware status indicators such as overall system status, chassis status, power supply, voltage, amperage status, etc.; Measures of the thermal status of the hardware including the current temperature, and the temperature status; Key metrics revealing the health of the cooling units/fans on the server, like, the current fan speed, and the fan status; Status of the power supply units

Siebel Enterprise	Siebel Web Server	<ul style="list-style-type: none"> Application access-related metrics such as, the number of times an application was accessed and for how long; Session activity on the web server measured by the number of current sessions to the server, the session duration, anonymous sessions removed/requested/returned, server response time, request time, etc.; Metrics revealing the lock behavior, which includes, the number and duration of user-initiated and anonymous locks
	Siebel Application Server	<ul style="list-style-type: none"> Metrics indicating the health of the object managers, such as, the current state of the object manager, and whether the object manager has reached its maximum tasks limit or not; Statistics revealing component health, including, the CPU time used up by components, the response time of the components, errors encountered by a component, etc.
	Siebel Gateway	<ul style="list-style-type: none"> Reports the errors logged in the gateway error logs; Indicates whether the Gateway Name Server is available or not
UPS Monitoring	APC UPS	<ul style="list-style-type: none"> Statistics revealing the health of the UPS batteries, which include, the current status (whether battery is low or normal) of the UPS batteries, the count of defective external batteries (if any), the power and voltage levels handled by the batteries, battery usage levels, whether the battery needs to be replaced, etc. Key metrics such as the UPS status, power and voltage I/O, etc.
	Generic UPS	<ul style="list-style-type: none"> Metrics revealing battery health, such as, the battery discharge time, battery runtime left, battery charge remaining, current battery status, etc.; Metrics that measure the quality of input and output lines, such as, the number of input lines utilized by UPS, the number of bad input lines to the UPS, the number of output lines utilized by UPS, etc.; Statistics revealing the voltage, current, and power inputs to the UPS via the input lines; Statistics revealing the voltage, current, and power output from the UPS via the output lines;
	XUpS	<ul style="list-style-type: none"> Metrics revealing battery health, such as, the battery runtime left, battery charge remaining, current battery status, etc.; Metrics reporting the utility line frequency, the number of times input was out of tolerance, source of input and output power, output load, etc.; The input voltage, input current, output voltage, output current, etc.

UPS Monitoring	Delta UPS	<ul style="list-style-type: none"> Statistics revealing the health of the UPS batteries, including the current status of the UPS batteries, the power, temperature and voltage levels handled by the batteries, the battery usage levels and whether the battery needs to be replaced etc, Key metrics revealing the number of low battery traps recorded, fuse failure events, power failure events, temperature failure events, overall failure events etc, Statistics revealing the voltage current and frequency inputs to the UPS via the input lines and the output lines; Key metrics revealing the UPS power capacity i.e., the load through each output line;
SAN Storage Devices	Hitachi AMS	<ul style="list-style-type: none"> Statistics revealing the health of each drive on the device, such as, the speed at which each drive processes read/write requests, etc. Statistics revealing the load on each drive on the device, such as, the rate of I/O operations on the drive, the rate of read and write operations to the drive, the transfer rate of read/write commands per drive, etc. Metrics indicating the control status of each port type on the device; Key metrics that shed light on the load on enabled ports, such as, the rate of I/O operations on each port, the rate of read and write operations per port, etc. Metrics that indicate the current mode of operation of the server; Metrics measuring the extent of usage of each processor supported by the device; Metrics revealing how effectively each cache on the device has been utilized, such as, the percentage of writes that are pending to every cache, the percentage usage of each cache's clean, middle, physical queues, etc. Metrics that point to the LUN on the storage device with the maximum throughput, such as, the rate of read/write commands executed on each LUN, the percentage of read and write requests that were served by this LUN, etc. Metrics indicating the overloaded RAID groups on the device, such as, the rate of read/write commands executed on each RAID group, the percentage of read and write requests that were served by this RAID group, etc.
	Clariion SAN	<ul style="list-style-type: none"> Statistics revealing the status, I/O operations and overall health of the storage processor, such as, status, percent utilization of storage processor, the speed at which each processor processes read/write requests, etc. Cache related metrics indicating the read/write hit ratios, the percentage dirty cache pages and the percentage dirty cache pages owned; Metrics that point to the LUN on the storage device with the maximum throughput, such as, the rate of read/write commands executed on each LUN, the percentage of read and write requests that were served by this LUN, etc. Statistics revealing the status, I/O operations and overall health of the disk, such as, LUNs count, percent disk utilization, Hard read/write errors, Soft read/write errors, read/write requests, etc.

		<ul style="list-style-type: none"> Statistics revealing the status, and overall health of the storage processor port, such as, status of link and port, the rate of read/write commands executed on each storage processor port, the number of data read/write commands made through this storage processor port etc. Metrics relating to the hardware status of the Disk Array enclosure such as the state of the Fans, Power state for each fan and the LCC state pertaining to each fan.
SAN Storage Devices	NexentaStor	<ul style="list-style-type: none"> Key metrics revealing how well each volume is capable of processing I/O requests; the current status of each volume and the metrics revealing the space utilization of each volume; Metrics revealing the physical memory utilization of the NexentaStor appliance; Users currently connected to the NexentaStor appliance;
	EMC VNX Storage	<ul style="list-style-type: none"> Measures revealing the space usage statistics of the File system such as the available space, total space, used space, percentage of space utilized etc., Metrics that point to the LUN on the storage system with the maximum throughput, such as, the rate of read/write commands executed on each LUN, the percentage of read and write requests that were served by this LUN, etc., Metrics revealing the I/O operations that are performed using the NFS protocol, the number of active threads in the NFS server etc., Statistics revealing the I/O operations that are performed using the CIFS protocol, the number of protocol connections and the number of files opened using the CIFS protocol etc. Metrics revealing the performance of the Data Movers by analyzing the number of running threads, joined threads, blocked threads etc., the idle/busy stat of the CPU, the available RAM size of the Data Mover etc., Metrics revealing the status of the statmon service, the connections that are availed and the percentage of connections that are used after the start of the statmon service etc., Metrics revealing the status of the Storage Port, the link state and the SFP state of the Storage Port. Statistics revealing the status, I/O operations, throughput and overall health of the storage processor, such as, status, percent utilization of storage processor, the rate at which each processor processes read/write requests, etc., Cache related metrics indicating the size of the read/write cache, the state of the read/write cache, the size of the free memory, the number of dirty cache pages and the cache pages owned, the physical memory of the storage processors etc., Statistics revealing the status, I/O operations and overall health of the disk, such as, LUNs count, percent disk utilization, Hard read/write errors, Soft read/write errors, read/write requests, etc., Metrics relating to the hardware status of the Disk Array enclosure such as the state of the Fans, Power state for each fan and the LCC state pertaining to each fan. Statistics revealing the network traffic in each network interface such as the I/O errors, packets/data transmitted and received etc.,

SAN Storage Devices	HP P2000 SAN	<ul style="list-style-type: none"> Measures revealing the status and health of the vdisk; Metrics revealing the I/O activity per vdisk, the rate at which read/write operations were performed, the rate at which data was read and written from the vdisk etc., Key metrics revealing the health and I/O activity per volume, the rate at which the read/write operations were performed, the rate at which data was read/written from the volume, the number of times the flush from the cache is/is not a full stripe Key metrics of the controller such as the percent of time the CPU of controllers were processing requests, the I/O activity performed on each controller, the rate at which data is transmitted, the rate at which data is serviced through read/write cache etc., Measures reporting the status and health of the expander port and host port of the controller; Metrics reporting the status, health and the failover status of the controller to the partner controller; Metrics revealing the I/O activity performed through each host port; Key metrics revealing the status, health and LED status of the disk; The I/O activity per disk , the rate at which data is transmitted etc., Hardware related measures such as the status of the expander, the SES elements of the expander and the sensors;
	Dell EqualLogic	<ul style="list-style-type: none"> External measures of the availability and responsiveness of the storage device; Measures revealing the number of disks and controllers inside the chassis that is to be monitored; Hardware-related metrics such as the speed and status of fans, the overall status of the hardware, the status of power supply units and temperature sensors, etc.; Key metrics pertaining to the space usage in the device, such as, total storage space, the space used, etc. Metrics revealing the errors and I/O activity per disk, such as disk state, disk size, number of disk errors, and the rate of bytes read/written; Metrics reporting the size and mode of the cache; Key space usage metrics pertaining to the storage pool, such as, the total space in the pool, the used space, reserved space, free space etc.; Metrics tracking iSCSI connections to the device, which include, the number of current connections, read/write latencies, rate of data transmissions/receptions, etc.; Metrics revealing controller health, such as, the status of the controller, the controller batteries, the temperature of processor and chipset, etc.; Critical metrics pertaining to the space usage in member arrays, which include, the total space in each array, the used space, etc.

SAN Storage Devices	Hitachi USP	<ul style="list-style-type: none"> Metrics revealing the current status of batteries used by each RAID store on the Hitachi USP storage device; Metrics that report the current status of the cache, controller, drive, fan, processor, power supply unit, shared memory used by each RAID store on the device; Metrics revealing the I/O activity for all the host bus adapters connected to each storage unit port, such as, the rate of I/O operations on each port, the responsiveness of the port to read-write requests, etc. Metrics revealing disk processor, channel processor, and DRS processor usage; Metrics revealing the I/O activity for all the host bus adapters connected to each storage unit port, such as, the rate of I/O operations on each port, the responsiveness of the port to read-write requests, etc. Measure indicating the percentage if data waiting to be written to the cache; Metrics revealing the I/O activity for all the host bus adapters connected to each storage unit port, such as, the rate of I/O operations on each port, the responsiveness of the port to read-write requests, etc. The percentage usage of the Cache switch to cache memory access path; Metrics revealing the I/O activity for all the host bus adapters connected to each storage unit port, such as, the rate of I/O operations on each port, the responsiveness of the port to read-write requests, etc. Metrics measuring the I/O activity on each logical volume, parity group, and LUN on the storage device such as, the rate of I/O operations on the component, the rate of data transfers on the component, the responsiveness of the component to I/O requests, etc.
	Hitachi VSP	<ul style="list-style-type: none"> The availability of the storage device over the network; The current status of critical hardware components such as the battery, cache, controller, drive, fan, processor, power supply and shared memory; Metrics revealing the I/O operations rate on each logical volume, the responsiveness of each volume to requests, percentage of data written to each logical volume, percentage of read requests served by each logical volume etc; Metrics revealing the I/O operations rate on each LUN, the responsiveness of each LUN to requests, percentage of data written to each LUN, percentage of read requests served by each LUN etc; Key metrics revealing the I/O operation rate of each port, the responsiveness of each port and the data traffic through each port; The utilization of the channels, disk and DRR processors; Do the caches have adequate memory space for storing data written to them, or are too many writes pending to the cache? The utilization of the cache memory to cache switch access paths;

SAN Storage Devices	Hitachi VSP Gx00 Series	<ul style="list-style-type: none"> • The availability of the storage device over the network; • The current status of critical hardware components such as the battery, cache, controller, drive, fan, processor, power supply and shared memory; • Metrics revealing the I/O operations rate on each logical volume, the responsiveness of each volume to requests, percentage of data written to each logical volume, percentage of read requests served by each logical volume etc; • Metrics revealing the I/O operations rate on each LUN, the responsiveness of each LUN to requests, percentage of data written to each LUN, percentage of read requests served by each LUN etc; • Key metrics revealing the I/O operation rate of each port, the responsiveness of each port and the data traffic through each port; • The utilization of the channels, disk and DRR processors; • Do the caches have adequate memory space for storing data written to them, or are too many writes pending to the cache? • The utilization of the cache memory to cache switch access paths;
	HP EVA StorageWorks	<ul style="list-style-type: none"> • Metrics revealing the current status of the temperature sensors, communication buses, EMUs, fans, modules, and power supply units within an array enclosure; • Metrics reporting the current status of each fan and power supply unit that is supported by the array controller; • Metrics revealing the I/O activity for all the host bus adapters connected to each storage unit port, such as, the rate of I/O operations on each port, the responsiveness of the port to read-write requests, etc. • Status and space usage related metrics pertaining to each disk group on the storage device; • Metrics revealing the level of traffic on each physical disk group, such as, the rate of read and write requests each group, the rate of data reads and writes to each group, etc. • Metrics revealing the I/O activity for all the host bus adapters connected to each storage unit port, such as, the rate of I/O operations on each port, the responsiveness of the port to read-write requests, etc. • The current state, formatting capacity, media accessibility, rate of disk read requests, rate of disk reads etc., of each physical disk on the device; • Metrics revealing how well the LUN cache services requests to each LUN group, such as, the rate of read and write requests to cache, time taken for reading from and writing to the cache, the time taken for reading from and writing to the physical disk, etc. • Metrics indicating the current operational state of each LUN on the storage device, and the current state of the read cache, write cache, and mirror cache of each LUN;

SAN Storage Devices	HP EVA StorageWorks	<ul style="list-style-type: none"> General health metrics such as the current operational state of the device ports on the controller, whether the device ports are up/down, the number and type of errors experienced by the device ports, the current health and operational state of the fibre channel ports, etc. Traffic-related metrics, such as, the rate at which the device received requests and transferred data; The current operational state of the EVA system, and the space usage by the system; Metrics revealing problem situations such as, cache battery failures, bad caches, bad mirror connections, bad mirror ports, controller failures, abnormal controller temperature, excessive CPU usage by controller, etc. Metrics that report the current condition and operational state of the host ports on the controller, the operational state of the external hosts connecting to the LUNs, the number of outstanding requests from external hosts, and the number of busy responses sent to the hosts, etc.; Metrics measuring the ability of the host ports to handle load, such as, the rate of read and write requests received by the host ports, the time taken by host ports for servicing read and write requests, etc.
	IBM DS RAID Storage	<ul style="list-style-type: none"> Metrics indicating the current status of the storage array and the space usage on the array; Metrics reporting the number and nature of errors experienced by the drive channels; Metrics indicating the current status of the drive channels and the drive channel links; Traffic-related metrics pertaining to the LUNs, such as, the number of I/O operations currently executing on the LUNs, the percentage of read and write operations, the cache hit percentage, etc. Metrics revealing the current status of the logical drives; Metrics revealing the status and extent of traffic on the controllers, which include, such as, the current status of the controllers, the number of I/O operations currently executing on the LUNs, the percentage of read and write operations, the cache hit percentage, etc. Metrics revealing the current status of the batteries, the fan canisters, the SFP transceivers, the temperature sensors, power supply units, fans in the array enclosure; Metrics indicating the current status of drive ports, host ports, and drives

SAN Storage Devices	IBM DS8000 Storage system	<ul style="list-style-type: none"> Metrics revealing the current health and operational status of each disk and the capacity of each disk; Metrics revealing the current health and operational status of each LUN, the I/O operations performed on each LUN such as the rate of read operations, write operations, percentage of read hits and write hits etc; Metrics revealing the current health and operational status of each FC port and the I/O activity on each FC port; Key metrics revealing the current health and operational status of each rank and the metric related to the space utilization of each rank; metrics revealing the I/O processing on each rank; Statistics revealing the current health and operational status of each array and the metrics revealing the space utilization of each array;
	NetApp USD	<ul style="list-style-type: none"> Measures that intercept SNMP traps and report error events that affect disk health, hardware failures, fatal errors that could cause system shutdown, warning events, etc. Status measures revealing the overall status of the NetApp Unified Storage system, the status of the AutoSupport feature, etc.; Measures revealing when and how many Consistency Points were triggered, such as, the total number of CPs, number of CPs that occurred due to a full NVRAM log, etc. Metrics revealing the percentage of time the CPU of the device was busy, the NVRAM DMA wait time, whether/not sufficient spare disks are available, etc. Disk usage metrics, such as, the RAID state of each disk, the total physical capacity of every disk, percentage of used space in each disk, read/write latencies per disk, etc. Metrics focusing on the performance of Aggregates, which include, the state of each aggregate, mirror state of an aggregate, the size of each aggregate, the percent usage of an aggregate, etc. RAID group usage measures, such as, number of prefailure disks per group, the total size of a group, the percent usage of a group, the percentage of media and parity scrubbing in a group, etc. Measures revealing the health of every Host Bus Adapter, such as, the current state of an HBA, whether/not this HBA is in standby mode, the queue depth of an HBA, total CRC errors, count of discarded frames and link breaks, etc. Key metrics related to current and potential hardware failures, which include, the current hardware temperature, the number of failed fans/power supplies, the current battery status, etc. Statistics related to block I/O protocols, such as, the availability status of each protocol service, the rate of I/O operations and I/O latencies experienced by each protocol, etc.

		<ul style="list-style-type: none"> • CIFS-related metrics, such as, the status of CIFS, the rate of read and write operations performed by CIFS, the latencies experienced by CIFS, etc. • Usage metrics pertaining to the WAFL buffer cache, such as, the rate of name cache, directory find, buffer hash, inode cache, and buffer cache hits and misses, the number of available buffers, blocks read from and written to the cache, etc.; • Metrics revealing iGroup configuration mismatches, such as, the count of igroups with invalid OS type, invalid user-partner settings, mismatched user-partner OS type setting, invalid ALUA setting, etc. • Measures revealing initiator status per iGroup, which include, the count of initiators with mismatching ALUA setting, OS type, VSA setting, and LUN mapping; • LUN-related metrics, such as, whether a LUN is online/offline, the space usage per LUN, the I/O operations and latencies per LUN
SAN Storage Devices	HP 3PAR Storage	<ul style="list-style-type: none"> • Key metrics revealing the current health and operational state of each cage, metrics revealing the I/O operations performed such as the rate of read operations, write operations, percentage of read hits and write hits etc; • Hardware related metrics such as the current state and health of each fan, running condition of each fan, health and operation states of each CPU; the current health and operational state of each power supply unit; • Metrics that reveal the current health and operational states of each controller and also points to the rate and percentage of read and write operations performed by each controller; • Statistics revealing the health and operational status of each disk, the speed at which each disk processes read/write requests, etc. • Key metrics that shed light on the load on enabled ports, such as the rate of I/O operations on each port, the rate of read and write operations per port etc.;
	Dell Compellent	<ul style="list-style-type: none"> • Statistics revealing the health and operational status of each disk on the storage center, the speed at which each disk processes read/write requests, etc. • Statistics revealing the rate of I/O operations on the storage array, the rate of read and write operations to the storage array, the transfer rate of read/write commands per storage array, etc. • Key metrics that shed light on the load on enabled ports, such as, the rate of I/O operations on each port, the rate of read and write operations per port, etc. • Metrics measuring the extent of usage of each controller in the storage center in terms of its read and write requests processing capabilities and also reveal operational and health state of each controller in the storage center; • LUN related statistics such as the current health, the rate of read/write commands executed on each LUN, the percentage of read and write requests that were served by this LUN, average time taken by each LUN to process I/O requests and number of requests performed by each LUN.

SAN Storage Devices	EMC XtremIO <ul style="list-style-type: none"> Statistics revealing the current state of each X-Brick, the number of SSDs and BBUs available in each X-Brick. Metrics that reveal the usage of CPU and state of each X-Env Statistics revealing the health and state of each SSD, space utilization, rate of I/O operations, bandwidth utilized by each SSD for performing read and write operations, the rate of read and write operations performed in each SSD, etc. Key metrics that shed light on bandwidth utilization of each data protection group, rate of I/O operations on each data protection group, the read latency per data protection group, etc. Metrics revealing startup time, health and connection state of the cluster, SSD space utilization by each cluster, rate of read and write operations, average read and write size, bandwidth utilized by each cluster, etc; Key metrics that shed light on health and current state of each target port, the rate of I/O operations on each target port , the rate of read and write operations per target port, etc. Metrics revealing number of LUN mappings to each volume, snapshots created by each volume, space utilized by each volume, etc. Measures indicating bandwidth utilization of each volume folder, rate of I/O operations on each volume folder, average read and write size of each volume folder, etc.; Statistics such as the space utilization of snapshots, number of volumes created using each snapshot, bandwidth utilized for performing read and write operations, rate of read and write operations of each snapshot, etc.; Metrics that point to the initiators with the maximum throughput, such as, the rate of read/write commands executed on each initiator, the number of read and write requests that were served by each initiator, bandwidth taken by each initiator to process read and write requests and read and write latency of each initiator. Number of sub-folders created from each initiator group folder, rate of read/write commands executed on each initiator group folder, bandwidth utilized by each initiator group folder to process read and write requests, etc; Measures revealing bandwidth utilization of each initiator group, rate of I/O operations on each initiator group, average read and write size of each initiator group, etc;
	Nimble Storage <ul style="list-style-type: none"> Measure indicating the number of times the events were triggered when failure of the disks, fans and the storage controller occurs; Hardware related metrics such as the failure of each power supply unit; temperature failure of the hardware components etc; Metrics revealing the read and write latency; the rate at which I/O operations were read/written sequentially/at random; the rate at which data is read from /written sequentially/at random.

		<ul style="list-style-type: none"> For each volume, key metrics that reveal the current state, space utilization and number of iSCSI connections etc; Cache related measures such as the rate at which read requests were successfully fulfilled by the read cache; Disk space utilized by the volumes and snapshots;
	Java Applications	<ul style="list-style-type: none"> Metrics reporting the number of classes loaded/unloaded from memory; JVM garbage collection related metrics, such as, the number of garbage collections started, the percentage of time the JVM spent on garbage collection, etc.; Metrics indicating the status of JVM threads, such as, the number of runnable, blocked, waiting, timed waiting, low CPU, medium CPU, high CPU threads, etc.; Metrics indicating the temperature and fan speed of the CPU supported by the local traffic manager; Resource usage metrics, such as, the CPU and memory usage of the JVM; Uptime statistics such as the total uptime of the JVM, the uptime during the last measurement period, etc.
Others	HP Blade Servers	<ul style="list-style-type: none"> Metrics indicating the current condition and composition (whether/not it has server blades, power supply units, temperature sensors, fans, fuses, net connectors) of the blade enclosure; Fan-related metrics such as the current condition of the fan, whether it is available or not, etc.; Fuse-related metrics such as the current status of each fuse, whether every fuse is available or not, etc.; Temperature sensor-related metrics such as the current condition and temperature of each sensor; Rack blade-related metrics, such as, the current status of each rack blade, the current condition, availability, power output, etc., of each power supply unit on each rack blade; Power enclosure-related metrics, that indicate the current availability, condition, and redundant state of each power enclosure; The type and current condition of each net connector
	IBM Storwize v7000	<ul style="list-style-type: none"> Measures indicating the I/O activity on the vDisk and the read/write latencies; the status of the vDisk, capacity, cache status of the vDisk; Measures relating to the data transmission and reception through each port, the rate at which the commands were initiated to the controllers and commands received through each port, link failures of FC port, numerical statistics of the FC port synchronization failure, signal loss of the FC port, invalid words transmitted through the port and invalid CRC etc., Measures relating to the data transmission and reception through each node, the latencies for each message excluding and including the time spent in the inbound/outbound queue;

	<ul style="list-style-type: none"> Measures indicating the time for which the CPU of the node canister was busy and the percentage of time the CPU was busy; Metrics indicating the rate at which the sectors were read and written on the vDisk; the rate at which pre stage sectors were initiated by the cache of the vDisk; the rate at which the sectors were written for track writes initiated by the vDisk cache etc., Measures revealing the status and capacity of the MDisk, I/O activity on the MDisk, the average time taken by the MDisk to respond to read/write requests etc., Measures revealing the status and capacity of the drive, the I/O activity on the drive, the average time taken by the drive to respond to read/write requests etc., Hardware related metrics such as the status of the enclosure, the numerical statistics of the total canisters in each enclosure, the online canisters, the power supply units, slots to accommodate drives in the enclosure etc., Hardware related measures such as the status of each port in the enclosure slot, status indicating whether a drive is present in the enclosure slot; the ID of the drive inserted in the enclosure slot; Metrics revealing the status of the power supply unit in the control enclosure, status of the canister; health of the battery; the charging status of the battery; current charge of the battery, end of life status of the battery etc.,
Others	Endeca Search <ul style="list-style-type: none"> Metrics revealing how well the Endeca Navigation engine is able to process search queries, which include, the number of successful search requests, the number of search queries waiting in queue, the average processing time of queries, the time spent in dgraph, etc.
	Bluecoat AntiVirus <ul style="list-style-type: none"> Metrics relating to the scan status such as the number of files that is scanned and the number of files that is detected; Key metrics relating to the status of resource usage and the percent of resource utilization.
	Cisco UCS Manager <ul style="list-style-type: none"> Measures revealing the overall health, operational state, performance state, the power, presence, thermal, and voltage states of core UCS components managed by the Cisco UCS manager such as the I/O modules, fans, chassis, fan modules, backplane ports, fabric ports, etc. . Measures revealing the overall health, operational state, performance state, the power, presence, thermal, and voltage states of the fabric interconnect PSUs Measures reporting the overall health, operational state, network load on each fabric interconnect ethernet port; Measures reporting the overall health, operational state, network load on each fabric interconnect FC port; Availability, operational state, and current configuration of the blade servers in each chassis managed by the Cisco UCS manager;

		<ul style="list-style-type: none"> Key metrics revealing the current temperature and input current of processors; Useful metrics indicating the power consumption and temperature of the motherboard on every blade server; Metrics revealing the overall health, performance, and load on each NIC managed by the UCS manager.
Others	Egenera PAN Manager	<ul style="list-style-type: none"> Measures revealing the current status of a PAN Domain, total number of pNodes available in the domain, the number of active pNodes, the percentage of pNodes in the domain that are currently in use, etc. Metrics related to PAN OPServers, which include, the current status of each PAN OPServer and the master/slave status of the PAN OPServer; License usage metrics, such as, the number of licences of each type available in the PAN Manager, the number of licenses of each type utilized, etc.; Measure revealing chassis health, such as, the current operating status of a chassis, type of switch available in a chassis, the number of uplink and downlink ports in a chassis, etc.; vSwitch-related metrics, such as, the vSwitch type, uplink status of a vSwitch, link dependency status of a vSwitch, the number of LPANs and pServers attached to a vSwitch, etc.; Statistics that measure pNode performance, such as, the current state of each pNode, the clock speed of the CPU in every pNode, the number of CPU cores available in each pNode, average CPU/memory/blade/hard drive/zone temperature of each pNode, etc.; Statistics revealing pServer performance, such as, the current operating state of each pServer, whether/not the PAN Agent is available on a pServer, the number of primary/failover pNodes per pServer, CPU/memory usage of a pServer, disk I/O of a pServer, network traffic handled by each pServer, etc.; Measures revealing the current status of each LPAN in the PAN Manager, the average percentage of CPU that is currently utilized by an LPAN, the number of pServers available in an LPAN, the number of pServers that are currently booting in an LPAN, etc.; Disk space usage metrics per disk partition per pServer, such as, the total capacity of a disk partition, the amount of available space per partition, the percentage of space usage on each disk partition, etc.
	2X Publishing Agent	<ul style="list-style-type: none"> Statistics that measure the CPU and memory usage of applications deployed on the server; Key metrics revealing whether connection to the gateway is available and the time taken to establish the connection;

Others	2X Terminal Server	<ul style="list-style-type: none"> • Redirector-related metrics, which include, the rate of data sent and received by the local server from the network, rate of network errors, reads/writes denied, etc. • Metrics tracking user profile size, such as, the current profile size, whether/not the profile exceeds quota, large files in user's profile, etc. • External metrics indicating availability of connection to server port and connection time; • Metrics measuring status of a user's login and the time taken to login; • Metrics revealing the resource usage of applications on the server; • Session-related metrics, which include, the number of idle/active/connected sessions to the server, etc. • Metrics revealing latencies in client connections to the server; • Metrics revealing the level of user activity on the server, which include, the current sessions for a particular user on the server, resource usage of the user, errors and bandwidth usage of the user, handles used by the user's processes, etc.
	2X Client Gateway	<ul style="list-style-type: none"> • Metrics indicating the availability of HTTP/HTTPPS connection to the gateway and the time taken to establish the connection; • Statistics that measure the health of the web server on which the gateway operates; • Metrics that capture the errors that occurred on the gateway; • Metrics that track the status of critical Windows services that support the client gateway; • Metrics tracking the usage of critical resources by the gateway processes;
	APP – V Management Server	<ul style="list-style-type: none"> • Measures revealing the availability and the authentication status of the APP-V Management server; • Measures revealing the availability and response time of the App-V database (generally a Microsoft SQL server), query execution time of the database server, the number of records fetched from the database server etc; • Measures revealing critical information relating to the App-V Publishing Admin log such as the information messages, critical messages, warming messages, verbose messages etc; • Measures revealing critical information relating to the App-V Publishing Operational log such as the information messages, critical messages, warming messages, verbose messages etc; • Measures revealing critical information relating to the App-V Client Management Admin log such as the information messages, critical messages, warming messages, verbose messages etc;

	App – V Client	<ul style="list-style-type: none"> Measures revealing the performance of different applications executing on the App-V Client; Measures relating to the App-V Publishing server such as the availability and the response time; Measures relating to the size of each application, the percentage of application that is currently loading, the usage of the application etc; Measures revealing critical information relating to the App-V Client Admin log such as the information messages, critical messages, warming messages, verbose messages etc; Measures revealing critical information relating to the App-V Client Operational log such as the information messages, critical messages, warming messages, verbose messages etc; Measures revealing critical information relating to the App-V Client Virtual Application log such as the information messages, critical messages, warming messages, verbose messages etc;
Others	TeraText ArborText	<ul style="list-style-type: none"> Measures revealing the number of rendering jobs that succeeded, failed, and are in queue;
	TeraText Content Server	<ul style="list-style-type: none"> External metrics of the availability and responsiveness of the Content Server; Measures revealing the file count and total size of each database; Disk cache usage metrics such as, the number and percentage of cache hits and misses, the percentage of disk cache in use, etc.; The memory usage of the Content Server; The number and duration of file operations and record-change operations; Measures revealing the time taken by the server to process security requests, Z39.50 operations, present operations, search operations, sort and merge requests;
	Marathon EverRun PVM	<ul style="list-style-type: none"> Metrics that reveal the current operational state, health state, and capacity of each disk partition available to the PVM and each disk partition of the XenServer host on which the PVM executes; Statistics reporting the current operational and health state of the XenServer master in every XenServer pool; Ethernet adapter-related metrics such as the current operational and health state of each Ethernet adapter of the XenServer host on which the PVM executes and every Ethernet adapter used by the PVM; Current operational and health state of each PVM

Others	Double Take	<ul style="list-style-type: none"> • Measure revealing the total memory allocated to the server; • Metrics relating to the uptime of the servers such as whether the server was rebooted or not, the time for which the server has been up since the last reboot, etc. • Metrics revealing the amount of time this connection has been active, current state, the number of operations currently in the retransmit queue on the source, the number of operations currently waiting in the acknowledgement queue, etc.; • Metrics relating to the security of the servers such as the number of successful logins, the number of failed logins, the number of login attempts that were successful during the last measurement period, and the number of login attempts that failed during the last measurement period.
	HP Enterprise Security Key Manager	<ul style="list-style-type: none"> • The utilization of each CPU on the HP ESKM appliance; • The memory utilization of the appliance; • Metrics indicating the total number of requests processed by the appliance; the number of successful requests and failed requests; • Key measures revealing the number of times power failure event/fan failure event/disk failure event was triggered on the appliance;
	Hitachi Compute Blade	<ul style="list-style-type: none"> • Key metrics revealing status related metrics of each blade server such as the current health, power supply, maintenance mode, etc; • The current voltage and power consumption of each blade server; • The current status of each LED and the color emitted by each LED that is available in each blade server as well as switch module; • Metrics revealing the current health, power supply, temperature, voltage and power consumption of each chassis; • Key metrics revealing the power status, operating mode and maintenance mode of each management module; • Metrics revealing the current status of each LED, color of each LED available in each fan module, the speed of each fan in each fan module; • The current health status, power status and redundancy status of each fan module; • The current health status and power supply status of each fan module;

	RSA Authentication Manager	<ul style="list-style-type: none"> Metrics indicating the total number of authentication requests processed by the appliance, the number of successful requests and failed requests and average time taken by the appliance to process authentication requests and responses; Measures revealing the number of authentications performed by the appliance, the number of successful authentications and failed authentications, and the number of new pin authentications and next token code authentications; Metrics revealing the size of the cache, the hit ratio of the cache and the number of times the cache was flushed; Key measures indicating the total number of authentication requests that were processed in each data source, the number of success requests and failed requests in each data source, the number of connections that are currently active in each data source, etc; The replication state of each replica instance; Measure reveals the number of sessions that are currently active on the appliance;
Others	Ruckus Zone Director	<ul style="list-style-type: none"> The CPU utilization of the Ruckus Zone Director; Memory related metrics such as total amount of the memory allocated, amount and percentage of utilized memory and the amount of memory that is available for use; Measures revealing the status of the Ethernet, the number of packets transmitted and received through the Ethernet, the amount of data received and transmitted through the Ethernet, etc; Key metrics revealing the unauthorized client devices related details such as the radio type, the channel, the rogue type, the encryption mode and the signal strength; Statistics revealing the total number of access points and the number of authorized devices and the unauthorized devices; Measures indicating the utilization percentage of the CPU and memory of WLAN; Metrics revealing the rate at which the number of packets were transmitted from and received by the Zone Director, the rate at which the amount of data were transmitted from and received, etc; Key measures revealing the rate at which the number of packets were transmitted and received by each user, the rate at which the amount of data were transmitted from and received, the number of packets that were dropped during transmission and reception of the packets, etc; Access point related details such as the total number of radios associated with each access point, the number of authorized devices associated with each access point, the number of multi cast and unicast packets that were transmitted and received, etc; Metrics revealing the resource utilization, the power management status of each access point radio, the number of packets transmitted and received, the number of successful authentications, the number of failed authentications, the transmitted power of each radio, etc;

Network File Systems	NFS on Solaris Server	<ul style="list-style-type: none"> Measures revealing the total number of RPC calls received by the server from clients, the number of bad calls, the number of calls with a length shorter than the minimum length requirement of an RPC request, the number of duplicate RPC requests, etc.
	NFS on Solaris Client	<ul style="list-style-type: none"> Measures revealing the total number of RPC calls made by the client, the number of bad length calls, the number of calls that timed out, the number of call failures and the reasons for the same, etc. Availability and access time of remotely mounted NFS file systems;
	NFS on Linux Server	<ul style="list-style-type: none"> Measures revealing the total number of RPC calls received from clients to the NFS server, the number of corrupted RPC requests, the number/percentage of RPC call failures, the number/percentage of bad authentication requests, etc.;
	NFS on Linux Client	<ul style="list-style-type: none"> Measures revealing the total number of RPC calls received from clients to the NFS server, the total number of retransmitted RPC calls from clients, the total number of times authentication information had to be refreshed, etc.; Availability and access time of remotely mounted NFS file systems; Measures indicating the availability and space usage of every NFS-mounted directory;
Cloud Providers	Amazon EC2	<ul style="list-style-type: none"> External metrics of the availability of HTTP/HTTPS connection to the cloud and the responsiveness of the cloud; Metrics revealing the availability of availability zones, regions, the time taken to access the regions, etc.; Metrics revealing the availability of instances launched on the cloud and the time taken to access the instances; Metrics reporting the number of instances on the cloud, the number of powered-on, powered-off, added, removed instances; Resource usage metrics revealing the percentage of CPU resources utilized by an instance, the network traffic generated by an instance, etc.; Metrics revealing the uptime of instances; Metrics capturing the failure of an instance to be powered-on/off in a region and the time taken for the same;

Cloud Providers	Microsoft Azure	<ul style="list-style-type: none"> • Key metrics revealing whether the test managed to establish TCP connection with the cloud, whether the DNS server was able to respond successfully to the request made to it, etc., • Connection related metrics revealing number of successful connections to the database instance, number of connections that experienced deadlock on the database instance, number of connections that were throttled on the database instance, etc; • Measures indicating number of database instances that are currently powered-on and powered-off on the cloud and number of database instances that were added to and removed from the cloud. • Storage related measures reveal the availability of the primary and secondary storage regions on the storage account, amount of ingress and egress data in the storage, etc; • Key metrics reporting number of anonymous, Shared Access Signature and authentication requests that failed due to authorization errors and network errors. • Metrics revealing deployment status of the cloud service, amount of data from and written to the disk allocated to the cloud service, amount of incoming and outgoing traffic, etc; • Measures reporting maximum number of disks allocated the virtual machine, amount of memory configured, etc; • Metrics revealing runtime availability, compute mode and site mode of the website, number of HTTP 2xx , HTTP 3xx, HTTP 401, HTTP 402, HTTP 403, HTTP 404, HTTP 406 and HTTP 4xx encountered by the website, etc; • Measures indicating number of VMs that were registered, added and removed from the cloud; • Key statistics revealing number of requests that are made, queued and rejected to the cloud service and number of ISAPI extension requests received on the cloud service; • Measures revealing availability of the cloud and time taken by the cloud to respond to the client requests; • Metrics reporting maximum number of cores, the storage accounts, cloud services, virtual network sites and local network sites that allocated to the subscription and number of cores, storage accounts, cloud services, virtual network sites and local network sites that are currently utilized in the subscription;
------------------------	------------------------	--

	Atlantis ILIO	<ul style="list-style-type: none"> • The comparative measures of the disk utilization with/without the Atlantis ILIO server, the total capacity saved in the target environment with the use of Atlantis ILIO server etc, • The I/O activity on the backend storage and the time taken by the Atlantis ILIO server to respond to requests; • The disk space utilized by the NFS datastore; • The read/write offloads by the Atlantis ILIO; • The I/O activity on the NFS datastore; • The NFS thread utilization of the Atlantis ILIO; etc.,
	Infoblox	<ul style="list-style-type: none"> • Metrics revealing the CPU and memory utilization, temperature of the hardware components of the Infoblox system; • The current status of each service running on the Infoblox system, the current status of each physical node, whether the Infoblox system is highly available or not; • Key statistics related to the messages that were transmitted/received through DHCP protocol and DHCP6 protocol; • Statistics revealing the number of DNS referrals in each zone, the responses that were successfully made to the appliance, the queries that failed etc, • Metrics revealing the rate at which queries were processed by the DNS cache, number of queries replied by an authoritative server during the last 5 minutes and last 15 minutes, number of queries replied by the authoritative server after referencing another server in the last 5 minutes and last 15 minutes etc;
	NetApp Cluster	<ul style="list-style-type: none"> • Key metrics revealing the percentage of time the CPU was busy performing system-level processing, metrics revealing the health of the hardware components of the NetApp Cluster such as the number of failed fans, power supplies, battery status etc; • Key metrics revealing the current state of each aggregate, the space utilization related metrics of each aggregate, the I/O activity of each aggregate etc, • Statistics revealing the current state of each disk and the I/O activity of each disk; • Statistics revealing the type pf jobs in the job queue and the count of jobs that were successful, running, failed, rescheduled etc, • Key metrics revealing the current state of the node and the I/O activity of each node, • The current state of each Vserver, the availability of the cluster peer, • The current state of each flash cache and current state of each Vserver peer relationship, • The overall health and state of each Fiber Channel Adapter, key metrics revealing the active iSCSI sessions, the processing capability of each session, the login failures, failed tasks and errors encountered by each iSCSI session;

	<ul style="list-style-type: none"> Key metrics revealing the processing ability of each logical interface, the errors encountered during data transmission/reception and the uptime of each logical interface; Statistics revealing how well data is handled by each FC port and the errors/failures encountered by each FC port; the availability of the FCP service; Statistics revealing the I/O operations performed using CICS protocol, the count of Windows/Unix user groups that crossed disk space and file usage quotas set etc; The current state, alignment state of each LUN configured in the NetApp Cluster, the space utilization of each LUN and the metrics related to the I/O activity of each LUN; The current state of each volume, the space utilization related metrics of each volume and the metrics related to the I/O activity of each volume etc; The current failover state of each cluster and the partner system;
Dell PowerEdge VRTX	<ul style="list-style-type: none"> Key metrics revealing the current state and health of each amperage probe, the input power of each probe, Hardware related metrics such as the current state and health of each fan, running condition of each fan, health of each cooling unit and DIMM; the current health of each PCI device; Hardware related metrics such as the current health of each power supply point, operational state of the power supply unit's sensor, health of each PSU, current health of each processor, battery and the overall health of the Dell PowerEdge VRTX, current health of each slot, current health of each voltage probe and temperature probe etc; The current health of each blade server, enclosure and the current status of the enclosure hardware; the current health of each RAID controller, etc Statistics revealing the current state of each physical disk and the metrics revealing the space utilization of each physical disk; the current health and operational state of each virtual disk etc;
IBM Integration Bus	<ul style="list-style-type: none"> Key metrics revealing the current load on the FTP server; Statistics revealing the number of connections available in the JDBC connection pools and the metrics revealing the utilization of the connections; the timed out JDBC connections; The number of times the statements were executed successfully for each ODBC DSN and the number of times the statements failed; the number of error prone connections in the ODBC DSN; The throughput of the messages through the SOAP service; faulty replies sent through the SOAP service; Key statistics revealing the number of connections handled by the firewall service,

	<ul style="list-style-type: none"> the number of open connections on the TCP client node and TCP server node; Statistics revealing the data and message transmission/reception on each TCP client node and TCP server node; Metrics revealing the number of requests to the CICS transaction server that were successful and the requests that actually failed; the number of decisions that were processed successfully by the Decision service; Key metrics revealing the garbage collection activity on the heap; time taken by the message node to process the input messages; the number of messages processed by each message flow node and the time taken to process the messages; Statistics revealing the number of messages processed by each message flow and the time taken to process each message in the message flow; The processing rate of each message in a thread and the number of input messages processed by each thread etc;
Real User Monitor	<ul style="list-style-type: none"> For each website, key metrics revealing the total number of page views that were viewed, Apdex score, average time taken by the pages to load, number of unique users who are currently accessing the web site, percentage of page views with normal user experience, page views that are slow in loading and page views that have encountered JavaScript errors, number of times this web site accessed by the mobile phones, tablets and desktops, time related measures such as average page rendering time, etc. For each browser, statistics revealing the total number of page views that were viewed, Apdex score, average time taken by the pages to load, number of unique users who are currently accessing the browser, percentage of page views with normal user experience, page views that are slow in loading and page views that have encountered JavaScript errors, number of times the browser was accessed from the mobile phones, tablets and desktops, time related measures such as average DOM ready time, etc; Key metrics disclosing the total number of page views that were viewed using each device, Apdex score of each device, average time taken by the pages to load on the device, number of unique users who are currently accessing the device, percentage of page views with normal user experience, page views that are slow in loading and page views that have encountered JavaScript errors for each device, time related measures such as the average DOM processing time, etc. Monitors the total number of page views of each page type, Apdex score, average time taken by the pages of this type to load completely on the browser, number of unique users who are currently accessing each page type, percentage of page views with normal user experience, page views that are slow in loading and page views that have encountered JavaScript errors for each page type, time related measures such as the average DOM processing time, etc.

	<ul style="list-style-type: none"> • Measures revealing total number of page views of each page group, Apdex score of each page group, average time taken by the pages of each group, number of unique users who are currently accessing pages of each group, percentage of page views with normal user experience, page views that are slow in loading and page views that have encountered JavaScript errors for each page group, time related measures such as the average TCP connection time, etc.; • Measures revealing total number of times pages were viewed by users from each region, Apdex score web site/web application based on the experience of users from this region, average time taken by the pages accessed by the users from this region, number of unique users who are currently accessing pages in this region, percentage of page views with normal user experience, the average front end time in this region, number of satisfied, tolerating and frustrated page views viewed by the users from this region, average DOM download time, etc • Metrics shed light on Apdex score of web site/web application based on the experience of users from this country, average time taken by the pages accessed by the users from this country, number of unique users who are currently accessing pages in this country, percentage of page views with normal user experience, the average front end time in this country, number of satisfied, tolerating and frustrated page views viewed by the users from this country, average DOM download time, etc.; • Statistics indicating Apdex score of web site/web application based on the experience of users from each city, average time taken by the pages accessed by the users from each city, number of unique users who are currently accessing pages in the city, percentage of page views with normal user experience, the average front end time in the country, number of satisfied, tolerating and frustrated page views viewed by the users from the city, average DOM download time, etc.;
--	--