



805 Veterans Blvd, Suite 309  
Redwood City, CA 94063

T (650) 260-8638  
E [support@daric.com](mailto:support@daric.com)

<http://www.daric.com>

## Software Performance, Scalability, and Availability Specifications V 3.0

Release Date: November 17, 2015

## Availability:

Daric guarantees a Monthly Uptime Percentage (defined below) of at least 99.95%, in each case during any monthly billing cycle. In the event Daric does not meet this guarantee, you will be eligible to receive a Service Credit as described below.

### Definitions

- “Monthly Uptime Percentage” is calculated by subtracting from 100% the percentage of minutes during the month in which any part of the Daric system, as deployed, was in the state of “Unavailable.”
- A “Service Credit” is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

### Service Commitments and Service Credits

Service Credits are calculated as a percentage of the total charges paid by you affected for the monthly billing cycle in which the service Unavailability occurred in accordance with the schedule below.

#### Monthly Uptime Percentage

#### Service Credit Percentage

Less than 99.95% but equal to or greater than 99.0%

10%

Less than 99.0%

30%

## Scalability:

Daric’s systems have processed approximately 125 million system transactions to date, with individual deployments seeing up to 500,000 customer-decision requests in a

given day. Automatic scaling allows for deployments to lower and raise capacity at any given time.

As Daric applications add more customer accounts and add feature functionality, there's a need to ensure that system performance keeps up pace with added scale to avoid performance bottlenecks. As there is a performance tradeoff between horizontal (scale out) and vertical (scale up) architecture, our systems are architected and built on a horizontal architecture.

Daric's core platform has been architected from the ground up to organically scale to massive volumes. The core system architecture and application are designed to scale horizontally at the application level, messaging infrastructure level as well as the database level. We use the following resources to ensure that the platform is scalable at all levels:

- Juniper & Mellanox for network gear

- SuperMicro high-capacity web and application servers

- FIPS compliant SafeNet encryption appliances

- High Performance SAN for storage layer

Load Balancer/Web Servers:

- All AWS and private web servers are horizontally scalable, with instances added on a per-need basis, and are configured in a redundant manner for HA

- Software load balancers route traffic to multiple servers

Application Servers:

Separate application servers for UI, synchronous and asynchronous transactions. All application servers are horizontally scalable and configured in a redundant manner for High Availability

UI and API processing is stateless, multi-threaded and configured and tuned with large heaps

Asynchronous transactions are re-entrant, highly multi-threaded and also configured with large heaps. They also segregate higher priority from lower priority processing

#### Messaging Infrastructure:

AMQ Broker used to implement publish/subscribe model

Broker is set up in active-passive configuration for quick failover/HA

#### Database Servers:

MySQL database. Sharding enables horizontal scalability at the database tier, with a master/slave setup for HA, with read-only slaves used to offload query workloads

AMQ database is used for persistent messaging in the event of failover recovery

#### Disaster Recovery

Sub-second data replication across all AWS facilities

Hourly backups for a revolving 24 hour window stored locally and offsite



805 Veterans Blvd, Suite 309  
Redwood City, CA 94063

T (650) 260-8638  
E [support@daric.com](mailto:support@daric.com)

<http://www.daric.com>

Daily backups for a revolving 30 day window stored locally and offsite

Weekly backups for a revolving annual window stored locally and offsite

12 months of data stored onsite and offsite in a PCI compliant manner

Backups are tested on a regular cadence, often daily, at a minimum weekly

Mock failover exercises are run as part of the business continuity plan

In addition to our AWS-based deployments, Daric leverages relationships with private datacenter facilities to ensure the highest levels of security, performance, availability and DR failover in the case of special non-cloud based client deployments. Both of the datacenter facilities are 100% synchronized. The second datacenter runs in warm standby and is able to take over full service capacity. Each datacenter is located in separate disaster zones -- Virginia (Herndon) and San Jose -- and are used by notable customers with high security and availability requirements.

Network infrastructure and latency

Fully redundant tier 1 carriers

1GB North/South Traffic, 40GB East/West traffic by Infiniband

Fully redundant Juniper switching and routing gear

Web and application firewalls to defend attacks

Access to a blended internet backbone comprised of 10+ carriers

DR site has 2 network carriers



805 Veterans Blvd, Suite 309  
Redwood City, CA 94063

T (650) 260-8638  
E support@daric.com

<http://www.daric.com>

## Performance:

### **Server and Render Response Times**

Daric guarantees that at an average deployment rate of 1,000,000 transactions per day it shall ensure that performance shall not fall below the following level:

95% of all queries for enterprise customer transactions respond in 8 seconds or less, including infrastructure, excluding backends.

### **Measurement Points:**

The response times will be measured using HP LoadRunner (or similar tool) located behind the firewall and in front of the web servers. The timer will measure the time from the request for a page to when the last bit required to render the page is returned. Backend response times will be measured using the application server log files.

### **Concurrency, Workload, and Throughput**

Daric guarantees that performance throughput for transactions will meet or exceed the standard requirements for a given deployment. Testing has shown throughput matches expectations for an average of 16 transactions a second, with over 500,000 transactions in a given day.