

## Trustgrid SD-WAN Load Testing Results

Completed By: Alkami  
Report Date: 04/28/2021

### Introduction

This document is a report of the testing plan and overall results of performance testing completed during a customer pilot executed by the Alkami Performance Team. These are the results of testing designed to model the most common sources and patterns of data between the Alkami data environments and customers systems.

### Executive Summary

Based on the testing completed, there does not appear to be any performance issues that can be directly ascribed to the Trustgrid SD-WAN solution.

### Customer Data Centers

The Alkami platform is architected to integrate with our customers and third parties to provide a pleasant and snappy user experience. This is best achieved with a high-performing connection between the Alkami environment and the customer data centers.

### Common Configurations

The sections below outline the various configurations of Alkami customer services. They can be summarized as:

1. On-Site
2. Pass-Through
3. Hosted Third Party


#### On-Site

Some customers choose to host their data on their own servers in a data center managed by their staff. Below is a high-level diagram of the points of connection between the Alkami environments and an on-site customer data center.



#### Pass-Through

Some customers choose to host their data in a third-party data center and they own the routing/connection to the data center. Below is a high-level diagram of the points of connection between the Alkami environments and an on-site customer data center and the final third party vendor.



## Hosted Third Party

Some customers choose to have their data hosted in a third-party data center available over direct connections or publicly available endpoints. Below is a high-level diagram of the points of connection between the Alkami environments and a third party data center.



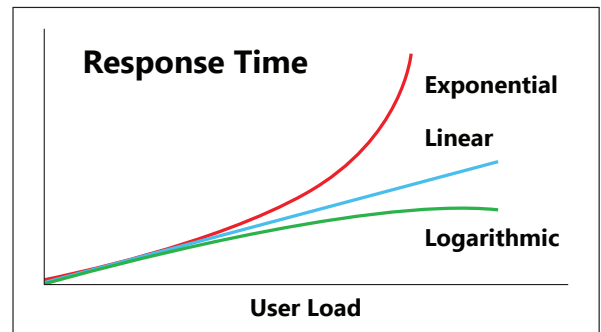
## Assumptions

The assumption for our testing is based on the idea that the Trustgrid solution is replacing the part of each diagram denoted with **"The Network"** with **Trustgrid**. This is analogous to replacing a networking cable with an equivalent cable based on the same specifications. We expect the performance to be roughly the equivalent to the current configuration performance. Though there are several kinds of connections diagrammed in the Common Configurations section, the pilot customers tested were selected, in part, because they represented **On-Site** Alkami customers.

## Measuring Performance

Performance is measured primarily based on response time for requests as compared to the throughput applied with the amount of data transmitted per request being consistently the same. Below is a picture of the relationship between response time and performance.

Our platform is working properly when we have Logarithmic performance. All customers tested represent systems that function with Logarithmic performance.



## Core Types

The Alkami Platform integrates with 3 basic kinds of cores, based on the network interactions and this required selecting and testing with a pilot customer that implemented each kind of service. For the purposes of this document, they will be designated as according to the table below:

Each customer was tested using the basic plan outlined in the Test Methodology section of this document.

Name	Service Type	Description
Web	Web/SOAP Based	HTTP/HTTPS connections that may or may not have connection pooling or re-use
OSI	Transient TCP Sockets	Many short-lived TCP connections
Sym	Persistent TCP Sockets	A few persistent TCP connections

## Test Methodology

### Creating Load

Since the Alkami platform is what will be generating the load in production, it makes the most sense to test using the Alkami platform rather than arbitrary or simplified data. The performance team has developed tools that control throughput using the parameter of user syncs per minute. Though users vary in complexity and data, the aggregate effect of calling the Alkami platform at a controlled rate will produce a flow of traffic that simulates the load expected in the real-world scenarios of the Alkami production environment.

### Scenarios

Testing was performed using 2 primary scenarios outlined below.

<b>Typical Load</b>	Data transmission for Pre-synced Users, which is the traffic we generate most of the time in production
<b>Heavy Load</b>	Data transmission for net-new Users, which is the traffic we generate during the process of onboarding a customer to the Alkami platform

The performance of both scenarios are calibrated to match the pilot customer with the corresponding core type. As an example, we might have:

Customer	Prod Throughput	Typical	Heavy
Web	100 Users / Min	100 Users / Min	50 Users / Min
OSI	75 Users / Min	75 Users / Min	50 Users / Min
Sym	90 Users / Min	90 Users / Min	50 Users / Min

The expectation is that **Typical** load is matched with the current customer load in production and that the **Heavy** load is consistent with our peak throughput, regardless of core type or customer throughput.

- **Baseline** tests were executed using the current connection used by Alkami to connect with the customer's systems
- **Comparison** tests were executed using the Trustgrid connection used by Alkami to connect with the customer's systems

## Reducing Variables

Our testing work is always about removing variables that could affect performance with the intention of using large samples of data with high duration and repeated execution. Below is a list of the variables we attempt to remove using our testing strategy.

### Code Version

For each customer tested we use the same

- Alkami Platform code version
- Load test environment
- Customer test core

### Test Sequence

Each test run has users that have met key requirements

- Curated users that have been confirmed as 100% valid on all systems
- Sequencing of users is the same, regardless of load levels
- Unique users are tested with no users repeated within a single run

### Good Testing

We ensure that each test is completed and meets our rigorous performance requirements for testing customers

- Key Performance Indicators - response times and throughput
- Success Rates - at least 97% test/user success rates
- Low error rates - no more than 3% for some services, even more strict for others
- Visual inspection of performance data to verify that tests executed properly for the duration of the test run.

### "Line Noise" or "The Internet"

- Testing is executed for a minimum of 30 minutes, if run as pairs or at least 45 minutes if a single run. This can be used to review if changes in performance are environmental or not.

## Results

Our testing results for these scenarios, by core type, are shown below.

### Web

The Web customer tests met all of the criteria for test quality outlined in the test methodology. The data below is the breakdown of metrics examined to compare baseline performance to Trustgrid performance at the various levels of load.

*Response times for all Web data are **displayed in seconds**, per scenario with the average across the sets of runs for grouping. The Grand Total line **averages** out the entire set of values across all runs for the specific connection type.*

## Core Response Time

The response times shown below are based on the call time to the MicroService, which implies the call times to the core. Our toolset, for this core type, did not enable measuring the actual wire times of the core calls, so we need to look for a trend on response times. The values are close to each other for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	50	0.451	0.456
	60	0.701	0.558
	70	0.661	0.547
History Total		0.604	0.520
Pre-Sync	70	0.482	0.470
	90	0.489	0.483
	110	0.442	0.520
Pre-Sync Total		0.471	0.491
<b>Grand Total</b>		<b>0.538</b>	<b>0.506</b>

## Bank Service - InitializeSyncStsId

The response times shown below are for the Bank Service call to IntiializeSyncStsId, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	50	0.641	0.640
	60	0.903	0.758
	70	0.863	0.751
History Total		0.802	0.716
Pre-Sync	70	0.663	0.653
	90	0.677	0.677
	110	0.625	0.710
Pre-Sync Total		0.655	0.680
<b>Grand Total</b>		<b>0.729</b>	<b>0.698</b>

## Radium Service - User Sync Job

The job execution times below are for the Radium User Sync Job, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	50	0.755	0.787
	60	0.990	0.827
	70	0.956	0.834
History Total		0.900	0.816
Pre-Sync	70	0.689	0.770
	90	0.701	0.703
	110	0.784	0.687
Pre-Sync Total		0.725	0.720
<b>Grand Total</b>		<b>0.813</b>	<b>0.768</b>

## Radium Service - Transactions Sync Job

The job execution times below are for the Radium User Sync Job, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	50	0.571	0.637
	60	2.661	2.202
	70	2.386	2.114
History Total		1.873	1.651
Pre-Sync	70	0.583	0.814
	90	0.598	0.616
	110	0.599	0.597
Pre-Sync Total		0.593	0.676
<b>Grand Total</b>		<b>1.233</b>	<b>1.163</b>

## OSI

The OSI customer tests met all of the criteria for test quality outlined in the test methodology. The data below is the breakdown of metrics examined to compare baseline performance to Trustgrid performance at the various levels of load.

*Response times for all OSI data are **displayed in seconds**, per scenario with the average across the sets of runs for grouping. The Grand Total line **averages** out the entire set of values across all runs for the specific connection type.*

## Core Response Time

Core response time had to be measured from two different sources, Bank and Radium. Response times are similar for both network solutions for all test scenarios.

## Bank Service

The response times shown below are for the Bank Service calls to the core. Our performance goals for this core are response times below 0.500 sec. These values are within our guidelines.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	20	0.147	0.089
	30	0.079	0.094
	40	0.084	0.088
History Total		0.103	0.090
Pre-Sync	20	0.081	0.088
	40	0.085	0.092
	60	0.080	0.087
Pre-Sync Total		0.082	0.089
<b>Grand Total</b>		<b>0.089</b>	<b>0.090</b>

### Radium Service

The response times shown below are for the Radium Service calls to the core. Our performance goals for this core are response times below 0.500 sec. These values are within our guidelines.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	20	0.260	0.181
	30	0.172	0.224
	40	0.175	0.178
History Total		0.202	0.194
Pre-Sync	20	0.078	0.083
	40	0.084	0.089
	60	0.077	0.094
Pre-Sync Total		0.080	0.089
<b>Grand Total</b>		<b>0.120</b>	<b>0.142</b>

### Bank Service - InitializeSyncStsId

The response times shown below are for the Bank Service call to IntiializeSyncStsId, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	20	0.724	0.619
	30	0.583	0.617
	40	0.583	0.590
History Total		0.630	0.609
Pre-Sync	20	0.641	0.609
	40	0.587	0.596
	60	0.557	0.584
Pre-Sync Total		0.609	0.596
<b>Grand Total</b>		<b>0.616</b>	<b>0.603</b>

### Radium Service - User Sync Job

The job execution times below are for the Radium User Sync Job, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	20	0.414	0.033
	30	0.319	0.343
	40	0.343	0.335
History Total		0.359	0.337
Pre-Sync	20	0.299	0.306
	40	0.310	0.319
	60	0.304	0.320
Pre-Sync Total		0.303	0.315
<b>Grand Total</b>		<b>0.322</b>	<b>0.326</b>

### Radium Service - Transactions Sync Job

The job execution times below are for the Radium Transactions Sync Job, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	20	3.556	3.139
	30	3.234	3.506
	40	3.208	3.234
History Total		3.333	3.293
Pre-Sync	20	0.242	0.246
	40	0.253	0.261
	60	0.237	0.564
Pre-Sync Total		0.245	0.357
<b>Grand Total</b>		<b>1.274</b>	<b>1.825</b>

### Symitar

The Symitar customer tests met all of the criteria for test quality outlined in the test methodology. The data below is the breakdown of metrics examined to compare baseline performance to Trustgrid performance at the various levels of load.

*Response times for all Symitar data are **displayed in seconds**, per scenario with the average across the sets of runs for grouping. The Grand Total line **averages** out the entire set of values across all runs for the specific connection type.*

### Core Response Time

The response times shown below are for the calls to the core. Our performance goals for this core are response times below 0.150 sec. These values are within our guidelines.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	50	0.088	0.093
	60	0.100	0.102
	70	0.117	0.120
History Total		0.102	0.105
Pre-Sync	40	0.039	0.091
	60	0.039	0.104
	80	0.039	0.115
Pre-Sync Total		0.039	0.115
<b>Grand Total</b>		<b>0.070</b>	<b>0.108</b>



### Bank Service - InitializeSyncStsId

The response times shown below are for the Bank Service call to IntiializeSyncStsId, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	20	0.736	0.789
	30	0.830	0.877
	40	1.000	1.045
History Total		0.855	0.904
Pre-Sync	20	0.576	0.734
	40	0.577	0.865
	60	0.583	1.530
Pre-Sync Total		0.579	1.043
<b>Grand Total</b>		<b>0.717</b>	<b>0.950</b>

### Radium Service - User Sync Job

The job execution times below are for the Radium User Sync Job, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	50	0.239	0.253
	60	0.267	0.280
	70	0.319	0.333
History Total		0.275	0.289
Pre-Sync	40	0.181	0.239
	60	0.179	0.278
	80	0.182	0.482
Pre-Sync Total		0.181	0.333
<b>Grand Total</b>		<b>0.228</b>	<b>0.303</b>

### Radium Service - Transactions Sync Job

The job execution times below are for the Radium Transactions Sync Job, which is one of our primary metrics for evaluating performance of the platform and the customer core.

The standards for performance set by Alkami are that this call should be executed within 2.5 sec. The values are well within the goals for average response time for all scenarios.

Average of CoreRT		Configuration	
Load Type	User Count	Baseline	Trustgrid
History	50	1.090	1.178
	60	1.243	1.294
	70	1.330	1.459
History Total		1.221	1.310
Pre-Sync	40	0.124	1.137
	60	0.129	1.279
	80	0.132	1.830
Pre-Sync Total		0.128	1.415
<b>Grand Total</b>		<b>0.675</b>	<b>1.345</b>