

PROJECT PROPOSAL

RADIUMBLOCK ENDPOINT SERVICE

	@radiumblock:matrix.org		@RadiumBlock		info@radiumblock.com
					+1-737-282-3611

[RadiumBlock](#) is a web3 infrastructure service provider, offering nodes to businesses and developers as a tool to help them build DApps faster, thereby eliminating the requirement for engineers to maintain and manage nodes. Using [RadiumBlock](#) services, blockchain aspirants can focus on their core business and strategies, relying on professionals to manage their blockchain infrastructure and performance.

About RadiumBlock

RadiumBlock is wholly owned by RealWeb LLC. RealWeb has been in business since 2004. Every year, our team manages millions of dollars in AWS infrastructure as well as managing our own server cluster deployment in data center locations across the world, making us one of the industry's most dependable and trustworthy web hosting/cloud services companies. Our engineers have extensive experience managing our own data centers and compute infrastructure. We also own and operate our own servers in multiple geography. We are now harnessing our expertise and experience in building reliable infrastructure services for the future of web3.

- We also provide validator services for [Kusama](#), [Polkadot](#), [Polkadex](#), [Sora](#), etc. We participate in the Polkadot and Kusama Thousand Validator Program, offering validator service to several other chains as well.
- We have recently launched our [snapshot](#) services (validator and active snapshots) which have been received well by the industry. In April 2022, RadiumBlock was awarded a grant from the Polkadot Treasury to build our [snapshot services](#).
- Earlier this year we launched the fastest (lowest latency), most geo-diverse(10 locations soon to be 15) endpoint service for the Kusama and Polkadot ecosystem. Decentralizing and geographically distributing the network is critical for ensuring that it can withstand any attack.
- We are also launching a Node-as-a-service platform to enable the DotSama ecosystem and parachains to launch validators, collators or full nodes.

Problems

Nodes are the backbone of blockchain technology. They keep track and manage all transactions on a network, while communicating with other nodes to confirm new blocks. The installation process for a node is not difficult. The lack of experience and technical challenges can make the job more complex than expected. And this is precisely the gap that RadiumBlock aims to address with their endpoint infrastructure.

Issues that you may face include but not limited to-

- **Performance**
 - o Latency Issues

	@radiumblock:matrix.org		@RadiumBlock		info@radiumblock.com		+1-737-282-3611
---	-------------------------	---	--------------	---	----------------------	---	-----------------

- o High Load
- **Complexity in**
 - o Scaling
 - o Fault tolerance
 - o Monitoring and maintenance
- **Costs** (Initial costs, Operating costs, etc)
 - o Time and money to handle complexity
 - o People needed

Scope of Work

This proposal is meant as an ongoing partnership between Astar and RadiumBlock. We aspire to apply our learning and insights in running Kusama/Polkadot endpoints and nodes to Astar infrastructure maintenance. Our endeavour is to become a trusted partner of Astar as we continue to grow as a highly robust RPC service providing company.

Objectives

To ensure that your system's performance remains high, it is important to –

- **Address Latency Issues**

With decentralized endpoints across all the continents in the world, RadiumBlock RPC services ensure that there is no latency irrespective of your geolocation.

- **Fault Tolerance**

RadiumBlock's security systems are state of the art, with an emergency response team on standby 24/7. Our strategically located data centres provide redundancy for fault tolerance and seamless connectivity for disruption tolerance. Our systems are configured to detect any downtime in any location. The node will be revived in another location with minimum disruption.

- **Cost-Effective**

Despite the highly advanced service setup, RadiumBlock's RPC services are highly cost-effective. Our systems are configured and optimized for the best

	@radiumblock:matrix.org	 @RadiumBlock	 info@radiumblock.com	 +1-737-282-3611
---	-------------------------	--	--	---

performance by scaling up whenever demand increases beyond the usual variability window and scaling down during off-peak hours.

Infrastructure Highlights

We offer a scalable, enterprise-grade infrastructure that can scale to handle traffic bursts and provide services in different regions around the world, while still delivering low latency responses. Additionally, the infrastructure is designed to be resilient and fault-tolerant irrespective of any downtimes in our service regions. The failover mechanism will patch over traffic from a failed data centre to other healthy data centres with minimal disruption.

To ensure a stable and optimal infrastructure, our team works around the clock using several monitoring tools.

The prominent features of our Endpoint services are:

- **Low Latency through Endpoint Delivery Network - What and Why it is Important**

- o Latency can make or break the experience of using blockchain technology. It is the measure of how quickly data can be sent from one point to another. The lower this number, the more responsive and real-time an experience will feel for users. Physics (the speed of light) won't allow us to get past the fact that there is physical distance involved which causes round trip time delays. Implementing an endpoint delivery network is the best way to dramatically decrease latency. With our distributed system, we can provide high-quality services with a low degree of latency. We are launching with 10 locations globally - Frankfurt, London, Mumbai, Ohio, Oregon, Sao Paulo, Seoul, Singapore, Stockholm and Sydney. Additionally, the company has plans on expanding into more countries and providing an even greater number of locations!

- **Monitoring and Observability -**

Our infrastructure expertise has enabled us to set a highly efficient monitoring system that works round the clock to ensure that deployed endpoints are up and running at all points in time. Any deviation is instantly handled with minimal disruption.

Some of the key areas of observability we focus on are:

- o *Endpoint availability and latency in different parts of the world*

	@radiumblock:matrix.org	 @RadiumBlock	 info@radiumblock.com	 +1-737-282-3611
---	-------------------------	--	--	---


We measure availability and latency from 28 points of observations around the world every 5 minutes. The system performs service monitoring by invoking a remote RPC call and an alarm is triggered if it takes more than 4 seconds to get a valid result.

Web Services - WebView 8 9 1
Updated: 04/05/2022 | 2:28:23 PM

Device Name	Targets	Status	Average	Frequency	Package	Last Response	Dallas	Denver	Miami	Minneapolis	New York	San Francisco	Seattle	Washington DC	Montreal
public wss endpoint for Kusama	1	OK	726 ms	5 min	Premium 10	2:23:24 PM	370 ms	272 ms	451 ms	415 ms	524 ms	305 ms	265 ms	650 ms	425 ms


o *Endpoint performance and load in each region of the world we deploy.*

We have dashboards that observe and report various system performance, load, and scaling parameters.



o *User experience of our Endpoint on the polkadot.js application*


We use a browser-based launch test that periodically loads our endpoint on one of the most demanded pages of the polkadot.js application (the staking page) - and observe from 6 different points in the world.



Device Name	Pages	Status	Average	Frequency	Package	Last Response	Dallas	New York	Copenhagen	London	Shanghai	Hong Kong
RB Kusama Public EndPoint Stakin...	1	✓	43.0 s	4 hours	Complimentary Green 5	2:27:46 PM	47.7 s	51.0 s	23.0 s	35.5 s	42.6 s	58.2 s

o Monitor the blockchain stats of the nodes.

Finally, we use our dedicated telemetry server to monitor the blockchain stats of all our endpoint machines across different geographies.



BEST BLOCK		FINALIZED BLOCK		AVERAGE TIME		LAST BLOCK	
#12,124,775		#12,124,772		5.997s		4.0s ago	

Endpoint	Location	Lat	Lon	Height	Block Hash	Time	Latency
kusana-endpoint	Stockholm	59.3293	18.0686	#12,124,775	0x7832d7f893298...	5.835s	0ms
kusana-endpoint	San Jose	37.3382	-122.2943	#12,124,775	0x7832d7f893298...	6.123s	157ms
kusana-endpoint	San Jose	37.3382	-122.2943	#12,124,775	0x7832d7f893298...	6.232s	239ms

• Auto-scaling and Load Balancing

Our team has implemented an innovative solution to help scale individual delivery points based on demand and balance the traffic across our computing infrastructure. We carried out a simulated stress test in order to define how much data can be handled by each compute unit before verifying if this system is capable of handling very large request loads.

• Fault Tolerance

It's not a matter of if the Internet will fail you, but when. Outages are a part of the game, and even the most well-designed Web 2.0 services can suffer from them sometimes!

We are very fortunate to live in an era where global infrastructure design allows for geographical as well as cognitive provider resilience. We have designed our global endpoint delivery network to be geographically diverse, which helps with the resilience of both geography and computing providers. We've successfully tested our infrastructure against "pull the plug" outages to see that the traffic is correctly and successfully routed to the endpoint in other parts of the network that are still functional.

- **Speed**

Radium Block's low latency and distributed delivery network spread across locations like Frankfurt, London, Dallas, Chicago, Seoul, Singapore, Sydney etc reduces time delay and provides seamless connectivity for disruption tolerance as can be seen in the following screenshots. The following table shows a comparison of the time taken for fetching blockchain status from different locations.

Monitoring Agent (Location)	RadiumBlock Duration	Others Duration
Sydney	248 ms	597 ms
Copenhagen	154 ms	655 ms
Shanghai	329 ms	438 ms
Tel-Aviv	290 ms	113 ms
Dallas	278 ms	588 ms
Frankfurt	183 ms	667 ms
Hong Kong	230 ms	470 ms
Miami	152 ms	546 ms
London	176 ms	628 ms
New York	343 ms	634 ms

Test conducted for Kusama on July 6th, 2022

- **Live Rolling Updates and Upgrades**

After years of experience in building cloud infrastructure, we know that software and other upgrades from time to time will be necessary. To keep the service running smoothly on a live network takes very delicate orchestration; but with Blue Green deployment pattern our team can design an upgrade strategy that ensures we can keep services online while upgrading or updating them as needed, all without disrupting end users' experiences. This paradigm is helpful even in an private deployment environment.

Decentralizing and geographically distributing the network is critical for ensuring that it can withstand any attack. Providing these endpoints requires significant resources which makes them difficult to achieve in a decentralized manner without compromising on other aspects like speed or security.

RadiumBlock's Future Plans

The Infura Problem:

The recent [Infura outage problem](#) seen in Ethereum was made worse by the temporary fork in the network it caused.

To tackle the forking problem, we are working on infrastructure solutions that can detect and notify users about any potential issues. Our goal is for this solution to eventually become a part of our regular monitoring framework and be made available for the community at large. We also believe that a diversity in endpoint providers helps detect and make the network more resilient to such problems.

Budget

Here is a budget breakdown per month assuming 5 geo locations on our private compute server clusters. This will allow for significantly more compute (approximately 20x vs AWS) and bandwidth (250TB at no cost vs ~\$31k in aws) to be provided to the endpoint service. The compute and bandwidth savings scales up as more geo locations are factored in. All of this explicitly tries to address the cloud cost over run and independence issue. Note, this is a budget breakdown in the interest of a public proposal. We can scale the location count to 10 or 15 based on preference.

Labor:

- Running DevOps, development to keep up with Astar (bug fixes, adjusting to new releases, security improvements, etc.):
 - o Fixed cost per 5 geo locations: \$ 1500/month

Operational Infrastructure costs:

Network	Type	Unit Price	Number of Units (per month)	Cost
Astar	Private Compute Cluster	\$500/mo	5	\$ 2500
	Load Balancer	\$ 200	1	\$ 200
	Storage	\$ 160/TB	5TB	\$ 800

	Log Storage	\$ 120/ TB	1TB	\$ 120
	Back-ups	\$ 0.058/GB-month	5TB	\$ 290
	DNS	\$ 200	1	\$ 200
	Traffic (Data Transfer)	\$ 0 upto 50 TB / month per geo	250TB	\$ 0
	Monitoring	\$50	1	\$ 50
	SubTotal			\$4160
	US Sales Tax			\$343
	Total			~\$4500

8.4 Total Budget:

Total for 5 geo-locations. This proposal covers our expenses for our infrastructure.

- Labor: \$ 1500/month
- Endpoint service, running on our private servers: \$4500/month
- *Total: \$6000/month*

We can also provide the following options:

10 geo locations: \$1200/month

15 geo locations: \$18000/month

We are assuming a uniform world wide distribution of locations.