



Office of Information Technology
Enterprise Infrastructure
Systems Engineering

Service Level Expectation (SLE) Virtual Infrastructure Hosting

Version 1.0

Effective: 11/01/2019

Document Owner:

Name	Position
Joni Loftus-Squillaro	Director, Enterprise Infrastructure
Steven Leonard	Manager of IT, Enterprise Infrastructure

Version History

Version	Date	Revision / Description	Author
0.1	10/10/2019	Original Draft	Andrew Page
0.2	10/15/2019	Revisions Added	Andrew Page
0.3	10/17/2019	Revisions Added	Steven Leonard
1.0	11/01/2019	Initial Version	Steven Leonard

OIT Enterprise Infrastructure Office of Information Technology

Thank you for partnering with us to help deliver IT services to the university community. This document is intended to set expectations about the service Enterprise Infrastructure Systems Engineering delivers as well as how to handle exceptions to that service.

Guiding Principals

1. The Enterprise Infrastructure Systems Engineering Group architects, deploys, operates, and maintains Enterprise-Class solutions in support of the shared hosting of University Services. Our infrastructure is designed to maximize and self-heal with technology in place to provide redundancy which allow availability and reduces the need for off-hours support.
2. For services that are architected to take advantage of the infrastructure, the Enterprise Infrastructure group will make all commercially reasonable efforts to restore service after a disruption.
3. Customer participation and availability is essential to the service restoration process.

Shared Responsibility Model

The Virtual Infrastructure operates under a shared responsibility model. Different components of a University Service are run by different groups. Enterprise Infrastructure (EI) offers three¹ different models for service. Customers unable to meet the stated customer responsibilities are encouraged to take advantage of OIT consulting services to fulfill the role.

¹ The Self-Managed Cluster option is in a Pre-release state. The Infrastructure is designed to support this option, but no customer has yet selected it.

² OS Management includes installation, and patching

	EI Managed Virtual Machine	Self Managed Virtual Machine	Self Managed Cluster
Application	Customer Managed	Customer Managed	Customer Managed
Custom Platforms	Customer Managed	Customer Managed	Customer Managed
Standard Platforms	EI Managed	Customer Managed	Customer Managed
OS Management²	EI Managed	Customer Managed	Customer Managed
Host Networking	EI Managed	Customer Managed	Customer Managed
VM Provisioning	EI Managed	EI Managed	Customer Managed
Hypervisor / Hardware	EI Managed	EI Managed	EI Managed
Storage / Networking	EI Managed	EI Managed	EI Managed
Data Center Facility	EI Managed	EI Managed	EI Managed
Utilities	Rutgers Facilities	Rutgers Facilities	Rutgers Facilities

Routine Maintenance

In order to maintain the efficiency and security of the Virtual Infrastructure, as well as to service customer requests, Enterprise Infrastructure preforms work according to the following Schedule:

Component	Risk	Announcement ³	Schedule
Network ⁴	Very Low	None	Work is done during the workday
	Low	None	Work is done from 6:00AM-8:00AM T, W, Th
	Medium-High	2 Weeks Prior	Work is scheduled between 4:00AM-6:00AM or by arrangement
OS Patching	Low	1 Week Prior	Machines are scheduled into groups so that Pre-production is patched before production and redundant machines are patched in different groups. Work is scheduled between 3:00AM-6:30AM M-F
	High	1 Month Prior	Work is scheduled over weekends or when least impactful.
Data Center Facilities	Low	1 Week Prior	Work is done during the workday
	High	1 Month Prior	Work is scheduled over weekends or when least impactful.

Work is announced on the “ei_maintenance” email list. New customers are required to create a mailing list for their department to self-manage local distribution. Customers are also expected to inform the users of their services as needed regarding any impacts to their services from routine maintenance.

Data Protection⁵

The Virtual Infrastructure includes the following Data Protections services

Component	Type	Frequency	Retention	Description
Block Storage (VDMK Files)	Full	2/Month	6 months	Primary Storage is RAID protected. Backups done via NetBackup and stored on DataDomain in the ASB Datacenter. Images are replicated off-site to Philadelphia
	Incremental	Daily	6 weeks	
File Storage (Isilon)	Daily	1/Day	31 Days	Frames are located in ASB and Knights Bridge Data Centers. Storage is replicated between frames. Snapshots are used for Data Protection
	Weekly	1/Week	6 Weeks	
Database	By Special Arrangement to ensure consistent backups.			

For systems that require a formal Disaster Recovery plan, Enterprise Infrastructure can work with Sungard Availability Services to provide a quote. This service is NOT included in standard pricing. The supporting business continuity plan is a responsibility of the department.

³ For an emergency or ongoing incident, changes will be applied without normal lead time.

⁴ The University schedules a Network Change Freeze during the beginning and end of each semester. Please see: <http://www.td.rutgers.edu/policies-guidelines/network-maintenance-freeze-periods/>

⁵ Backups can be customized by request, and those customizations will override these defaults.

Service Requests

A Service Request is defined as a request for information about or a change to the existing service. They may also be opened to report an incident.

To open a service request please visit <https://hdrti.rutgers.edu/ess/> and select “EI Virtual Managed Services” as the request category

Submit New Request

SUBMIT REQUEST

Request Category: **EI Virtual Managed Services** ▼

Request Type: ----- SELECT A REQUEST TYPE ----- ▼

Description:

You will receive a ticket number via email when opened. Email notice will again be sent when your issue is resolved. Often EI staff will require additional information in order to process your request. Service Requests that require an additional customer response that goes unanswered for 3 or more days may be closed with comment “3 days no client response”. The SR can be reopened as necessary per client request.

Service Priority

Priority	Description	Expected Acknowledgement Time.	
		Normal Hours	Off Hours
P1	Major Interruption to critical service(s) Failure of redundant Infrastructure Customer also commits resources 24x7 until resolved (Note: Services that are not architected for high availability do not qualify for P1 Status)	15 Minutes	15 Minutes – 1 hour
P2	Significant impact to service	15 Minutes	Best Effort
P3	Service degraded	1 Business Hour	N/A
P4	Routine Service Request	4 Business Hours	N/A

Normal Service Hours

- **Core EI Systems Staff Service/Support Hours:** 8:00am – 5:30pm Monday through Friday
Note: excluding University holidays and closures

Incident Communications

Enterprise Infrastructure communicates incidents via a web page at:

<https://status.ei.rutgers.edu>

Customers are encouraged to subscribe to incidents via that website. All Customers are responsible for crafting messages to their user communities.

Larger incidents which need to be communicated to the larger university community are handled through the OIT Enterprise Service Delivery Team.

Escalation Process⁶

The normal service request process should be used whenever possible. When a P1 or P2 Incident happens outside of Normal Service Hours, customers can begin our escalation process:

1. Please submit a normal service request. This step should only be skipped if the service request system is broken. Please include a way to get in contact off-hours.
2. Check our status page for a list of known impacts to the virtual hosting page. If a known issue is listed, your service request will be seen by staff working the issue **(Coming Soon)**
3. If you need to reach the on-call technician, please call +1 (224) 269-7273
LISTEN TO THE PROMPTS AND CHOOSE THE BEST OPTION
4. The system will attempt to connect you live to an EI staff member. After 60 seconds you will be prompted to leave a message. Please leave all proper contact information for an EI staff member to properly return your call. Leaving incomplete information could/will delay the escalation and resolution of service(s) or incident(s) being reported.

Once a call is placed an alert is created in our Opsgenie instance. This cloud-based tool will follow a pre-defined escalation policy to reach someone in the Enterprise Infrastructure Systems Engineering Group. The policy is as follows

1. The call is routed live to the on-call technician. If they can be reached, you will be connected via voice. The tech will ask for your contact information and ticket number. They will arrange for a call back.
2. If the call is not connected quickly, the system will take a voice mail. This voice mail will be attached to an alert which will follow the normal escalation path.
 - a. The on-call tech is given 5 minutes to acknowledge the alert.
 - b. The other members of the on-call rotation are contacted and given 5 minutes to acknowledge the alert
 - c. After 10 minutes, the alert is escalated to Enterprise Infrastructure Leadership team.
 - d. Alerts continue to repeat until acknowledged.

Please note, the timeframe given are for acknowledgement. All on-call technicians in Enterprise Infrastructure, are capable of basic triage for incidents, but may need to engage other teams during problem resolution. Only one staff member is officially on call. Reaching other resources will continue to be best effort.

⁶ Significant changes to the escalation process are expected as part of the IT Service Management / Service now project. Please make every effort to refer to current information before starting the process.

Escalation Flow

