elastic

Elastic Stack
最佳实践

Agenda

1. Elastic Stack 架构
2. Elasticsearch 最佳实践
3. 如何升级
4. 性能优化
5. 问与答

The Elastic Stack 7.0

Elasticsearch queries got faster, Kibana got a makeover, and that's just the start. Dive into features like faster top k queries, nanosecond timestamps, Function Score 2.0, and so much more.

Download Read More

Deploy Elastic Cloud Enterprise 2.2 with role-based access control and a UI for cross-cluster search.

Learn More

Try new production-ready features like cross-cluster replication in 6.7.

Learn More

Analyze your geospatial data with the new Maps app in Kibana.

Learn More

Elastic Stack —一体化的完整数据处理堆栈，从数据摄入到分析展示、价值获取

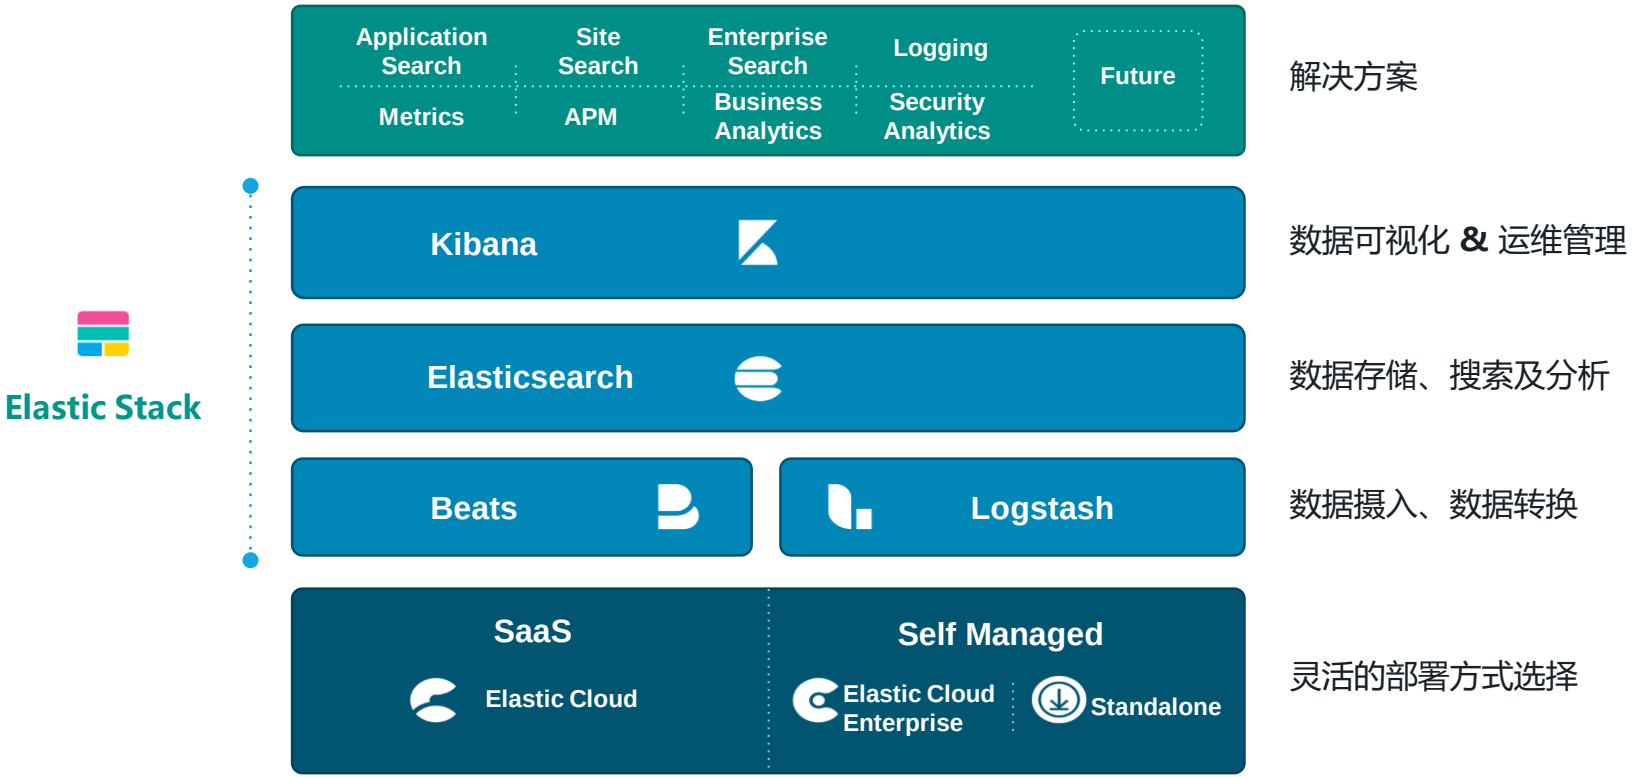

elastic

PoC 架构：上马快

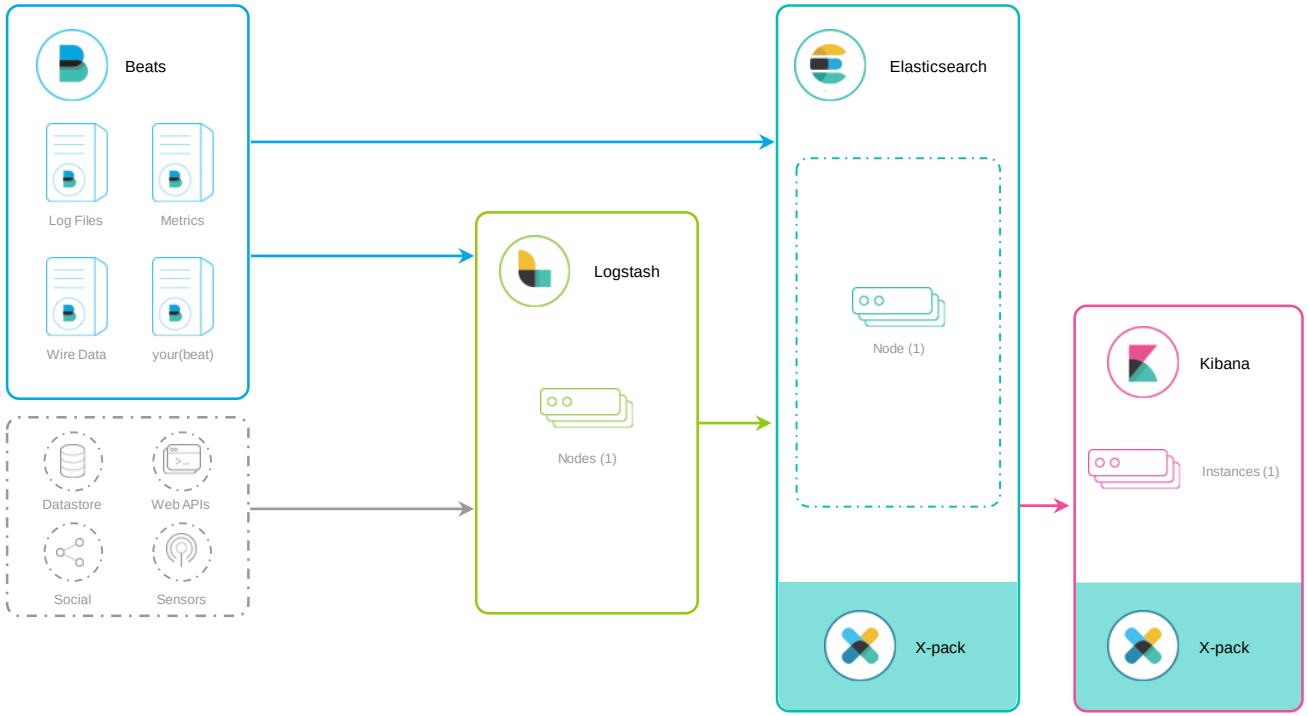

elastic

一周后

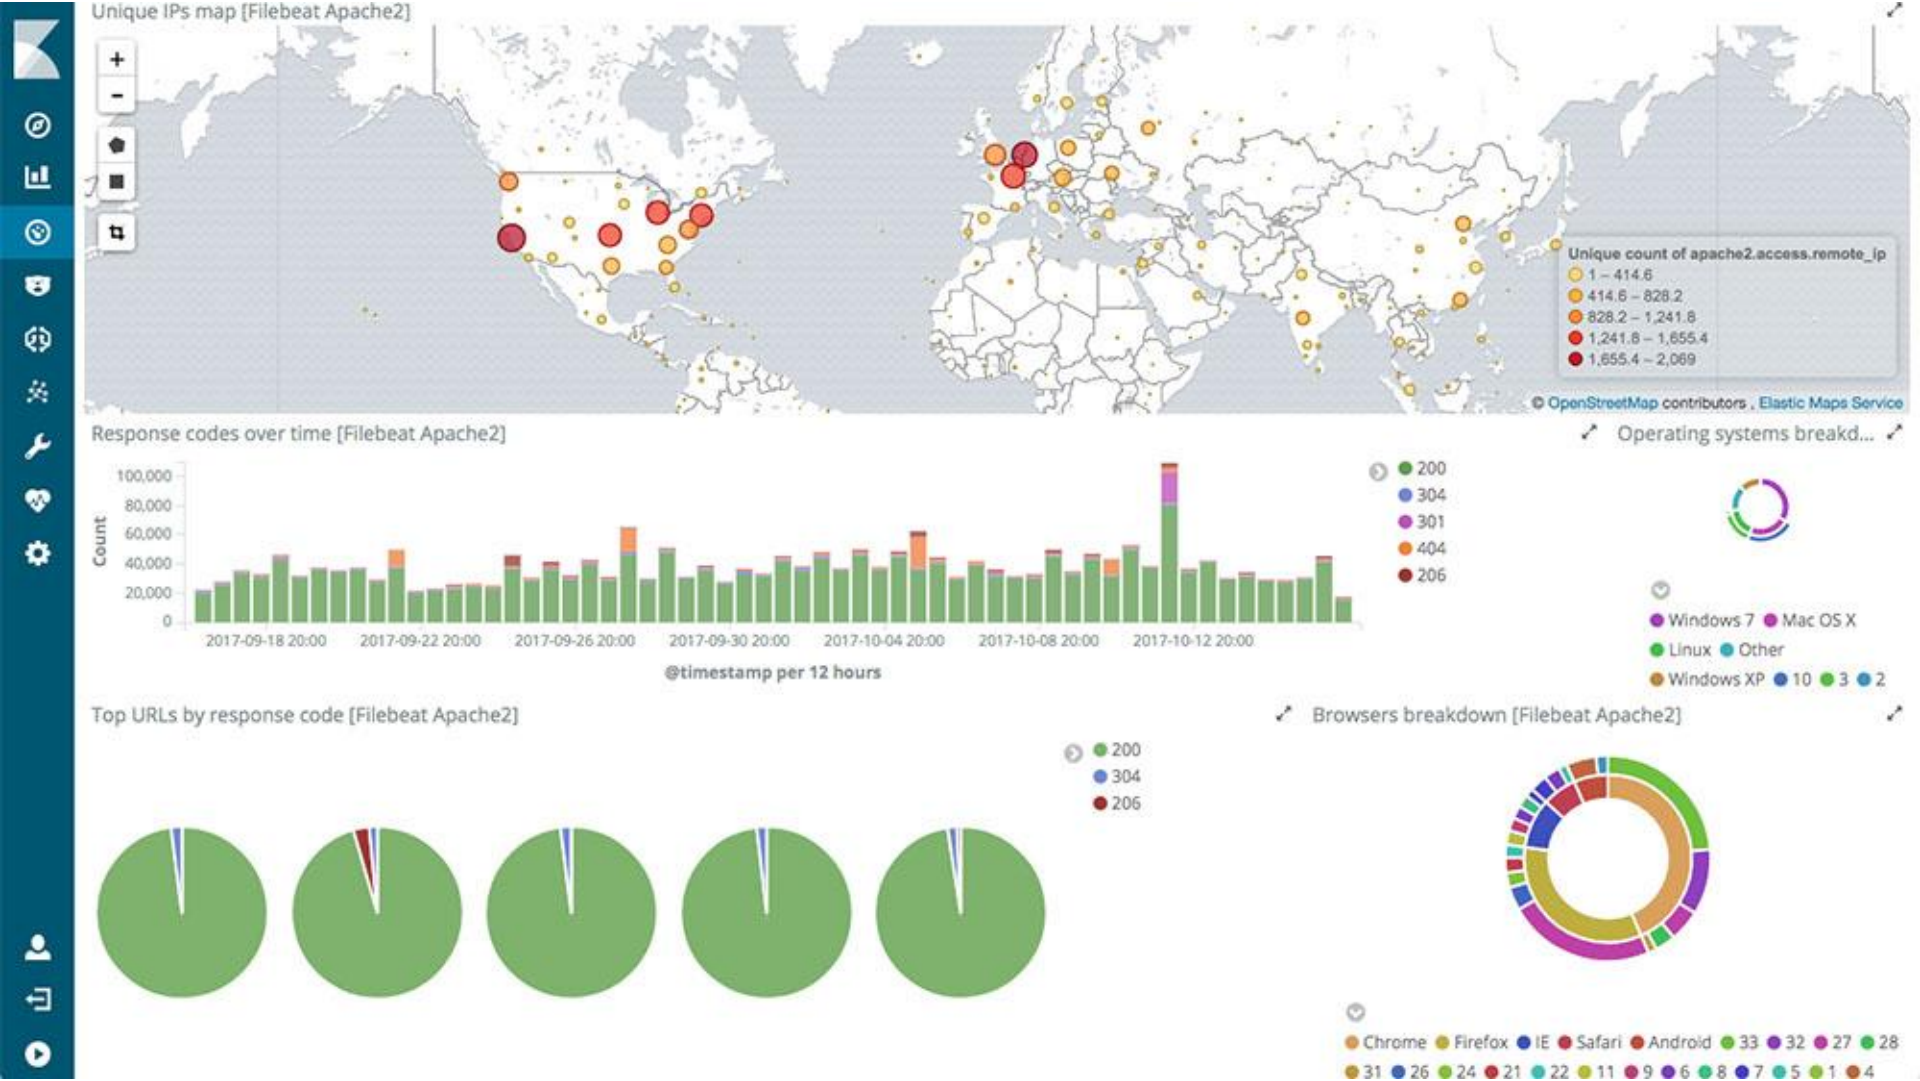
Unique IPs map [Filebeat Apache2]
Unique count of apache2.access.remote_ip
1 – 414.6
414.6 – 828.2
828.2 – 1,241.8
1,241.8 – 1,655.4
1,655.4 – 2,069
© OpenStreetMap contributors , Elastic Maps Service
Response codes over time [Filebeat Apache2]
Count
100,000
80,000
60,000
40,000
20,000
0
2017-09-18 20:00
2017-09-22 20:00
2017-09-26 20:00
2017-09-30 20:00
2017-10-04 20:00
2017-10-08 20:00
2017-10-12 20:00
@timestamp per 12 hours
200
304
301
404
206
Operating systems breakd...
Windows 7
Mac OS X
Linux
Other
Windows XP
10
3
2
Top URLs by response code [Filebeat Apache2]
200
304
206
Browsers breakdown [Filebeat Apache2]
Chrome
Firefox
IE
Safari
Android
33
32
27
28
31
26
24
21
22
11
9
6
8
7
5
1
4

*

Uses lucene query syntax

Add a filter +

Dashboards [Filebeat System]

Syslog | Sudo commands | SSH logins | New users and groups

Syslog events by hostname [Filebeat System]

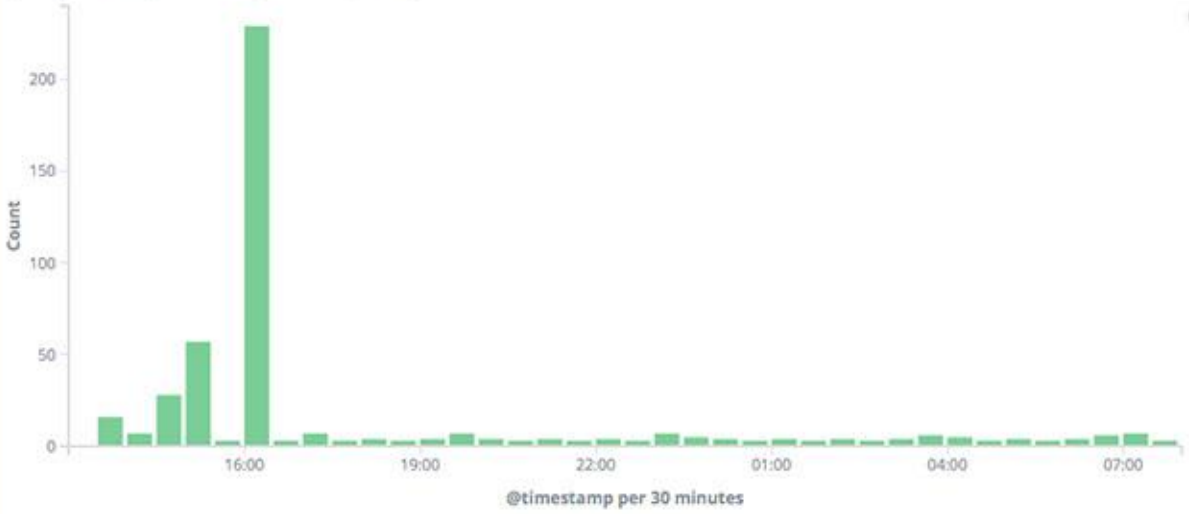

Syslog hostnames and processes [Filebeat System]

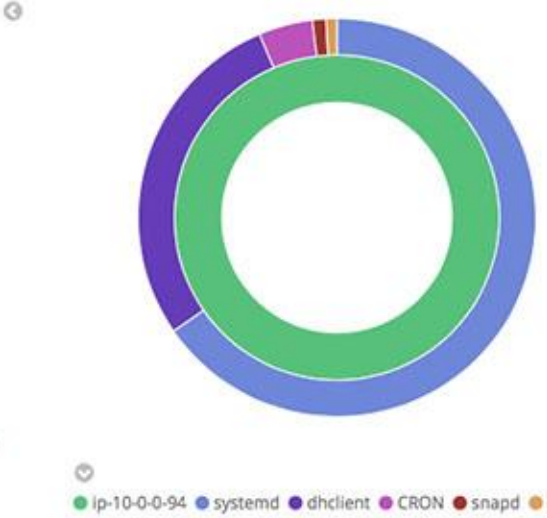

Syslog logs [Filebeat System]

1–50 of 470

Time	system.syslog.hostname	system.syslog.program	system.syslog.message
October 11th 2017, 07:50:05.000	ip-10-0-0-94	dhclient	DHCPREQUEST of 10.0.0.94 on ens3 to 10.0.0.1 port 67 (xid=0x3c54b2d6)
October 11th 2017, 07:50:05.000	ip-10-0-0-94	dhclient	DHCPACK of 10.0.0.94 from 10.0.0.1

*

Uses lucene query syntax

Add a filter +

Slow queries over time [Filebeat MySQL]

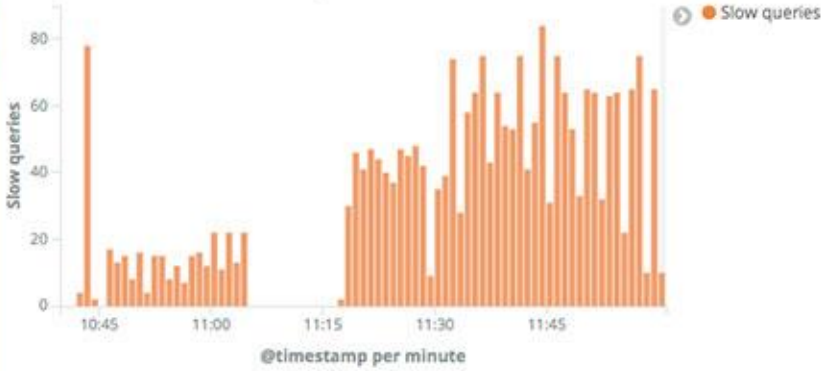

Error logs over time [Filebeat MySQL]

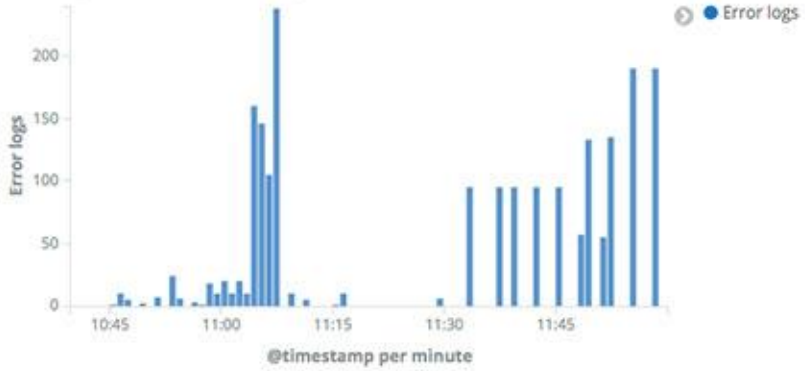

Slow logs breakdown [Filebeat MySQL]

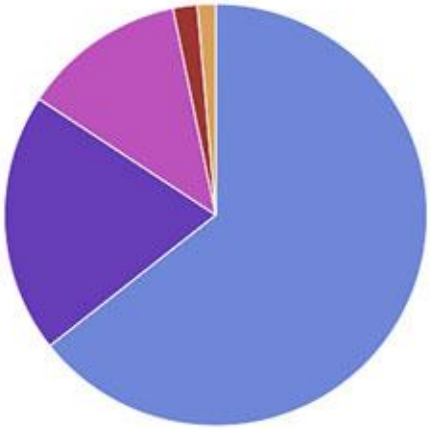

Error logs levels breakdown [Filebeat MySQL]

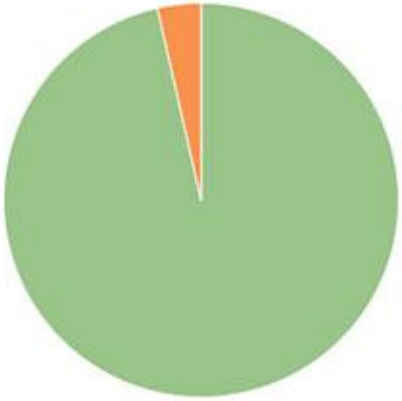

生产环境基本架构：扩展易

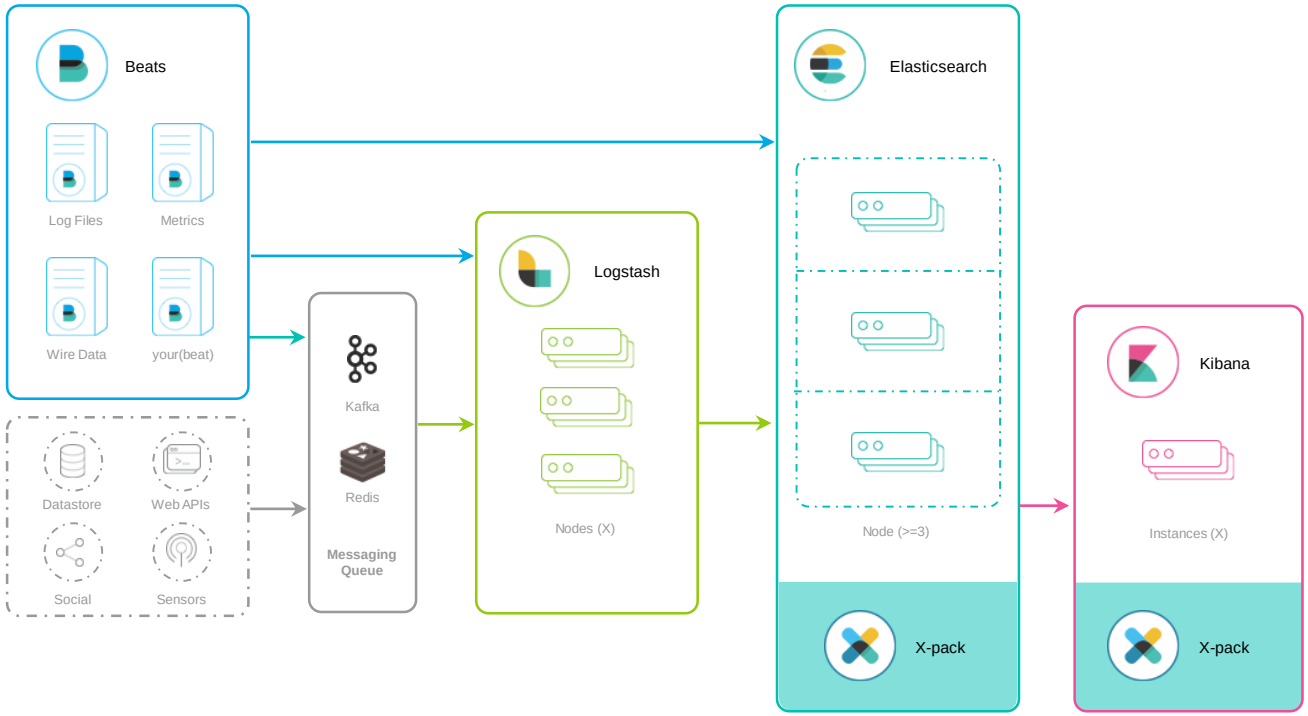

elastic

一个月后

两个月后

热数据与冷数据

热数据
当日数据：有读有写
最近1日数据：频繁读取

温数据
2日~7日数据：偶尔读取

冷数据
8-30日的数据：极少读取

elastic

热节点与冷节点

热节点
高配，内存磁盘比高
8Core+/64GB/2TB SSD

温节点
中配，内存磁盘比适中
8Core+/64GB/4TB SAS

冷节点
低配，内存磁盘比低
8Core+/64GB/6TB SAS

index.routing.allocation.require.box_type:hot

elastic

生产环境高级架构：冷热架构

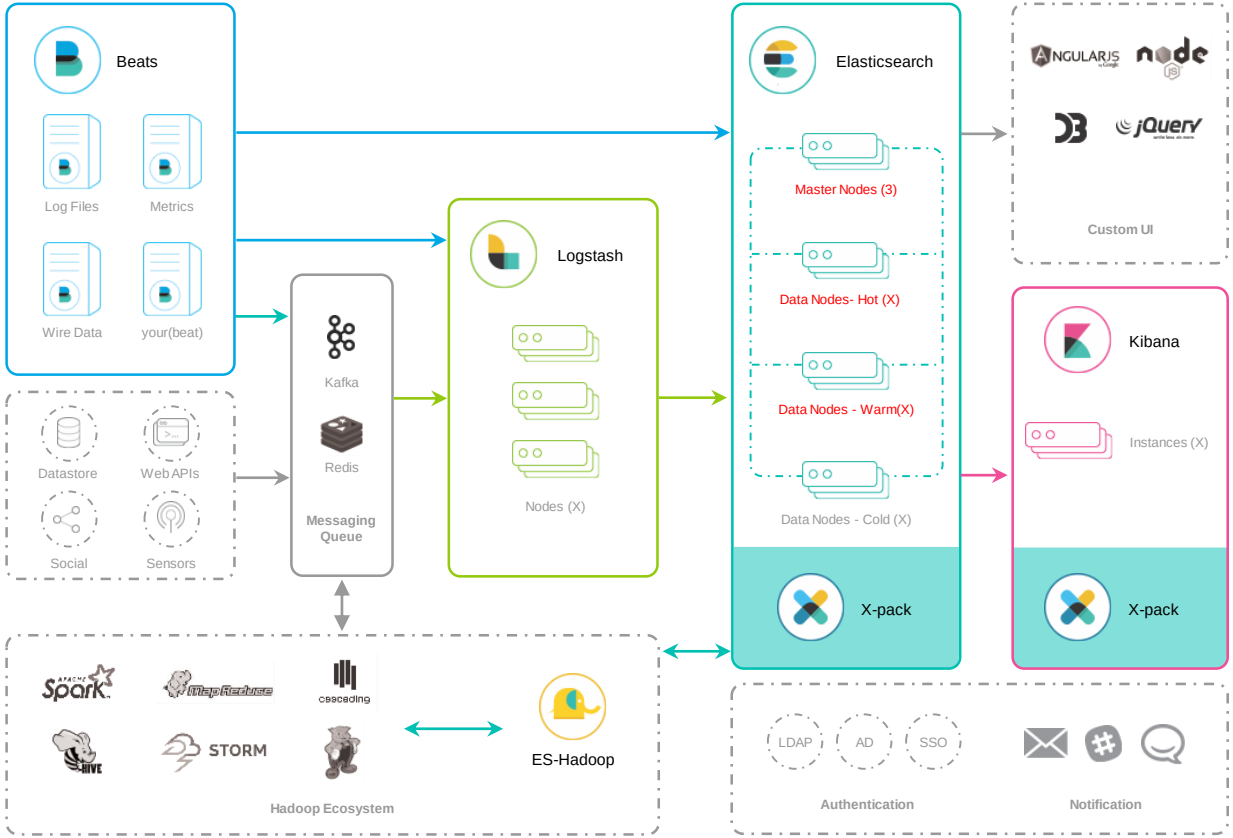

Agenda

Use color to highlight

1. Elastic Stack 架构
2. Elasticsearch 最佳实践
3. 如何升级
4. 性能优化
5. 问与答

三个月后

最佳实践

- 按时间建索引，每天自动删除过期的索引
- 使用 Index Template 模板，合理设计 Mapping
- 分片数不能过多也不能过少，分片大小控制在20GB 以内
- 使用别名
- 定期做 Force Merge
- 冷热数据分离，定时做迁移任务
- ……

Index Mapping

Index Alias

Shard Count Setting

Node Attributes

Curator

Replica Setting

知识点

Delete Index

Force Merge

Index Read-Only

Shard Allocation

Index Template

Create Index

Rollover

Shrink Index

ILM(Index Lifecycle Management) 最佳实践的集大成者

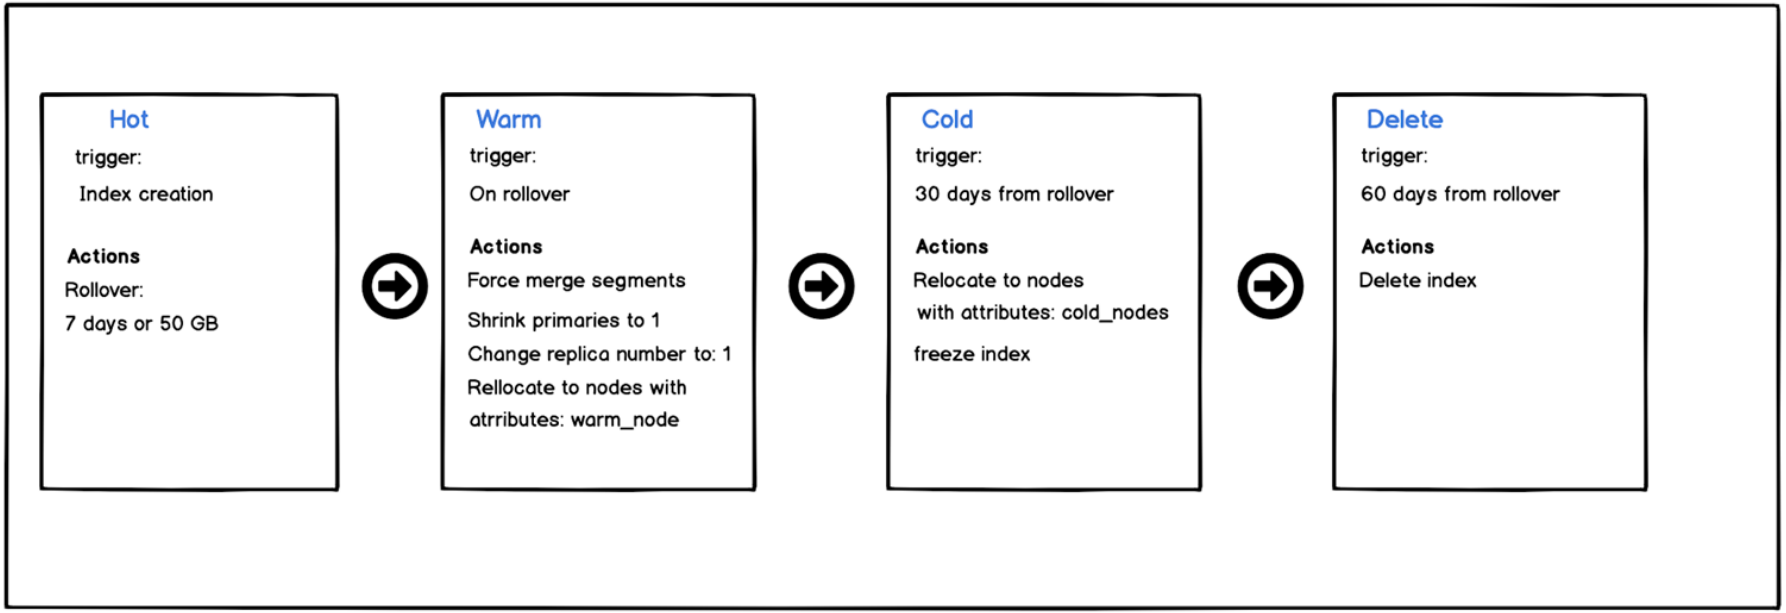

elastic

索引生命周期管理 Index Lifecycle Management

Hot Phase - Index to `my-logs`, Search on `my-logs`

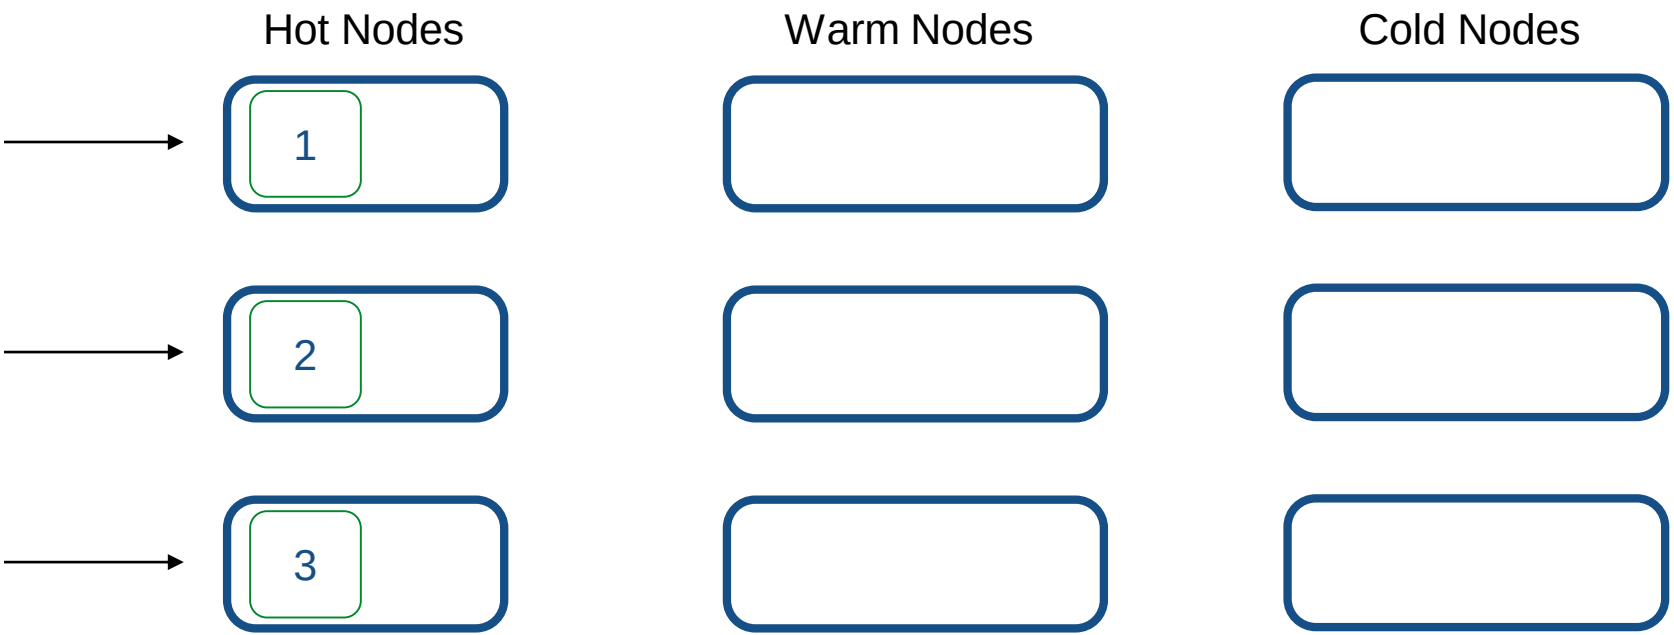

索引生命周期管理 Index Lifecycle Management

Hot Phase - Rollover

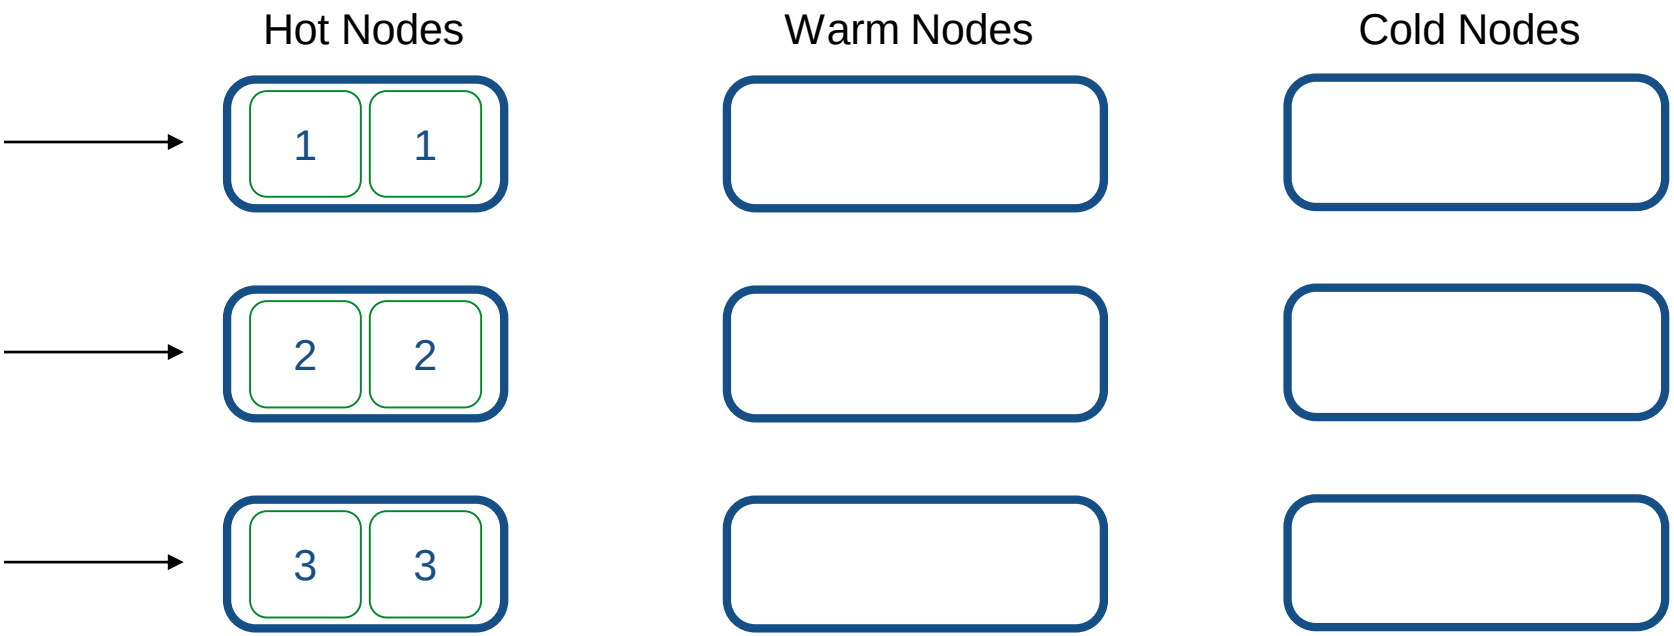

索引生命周期管理 Index Lifecycle Management

Warm Phase -

Allocate

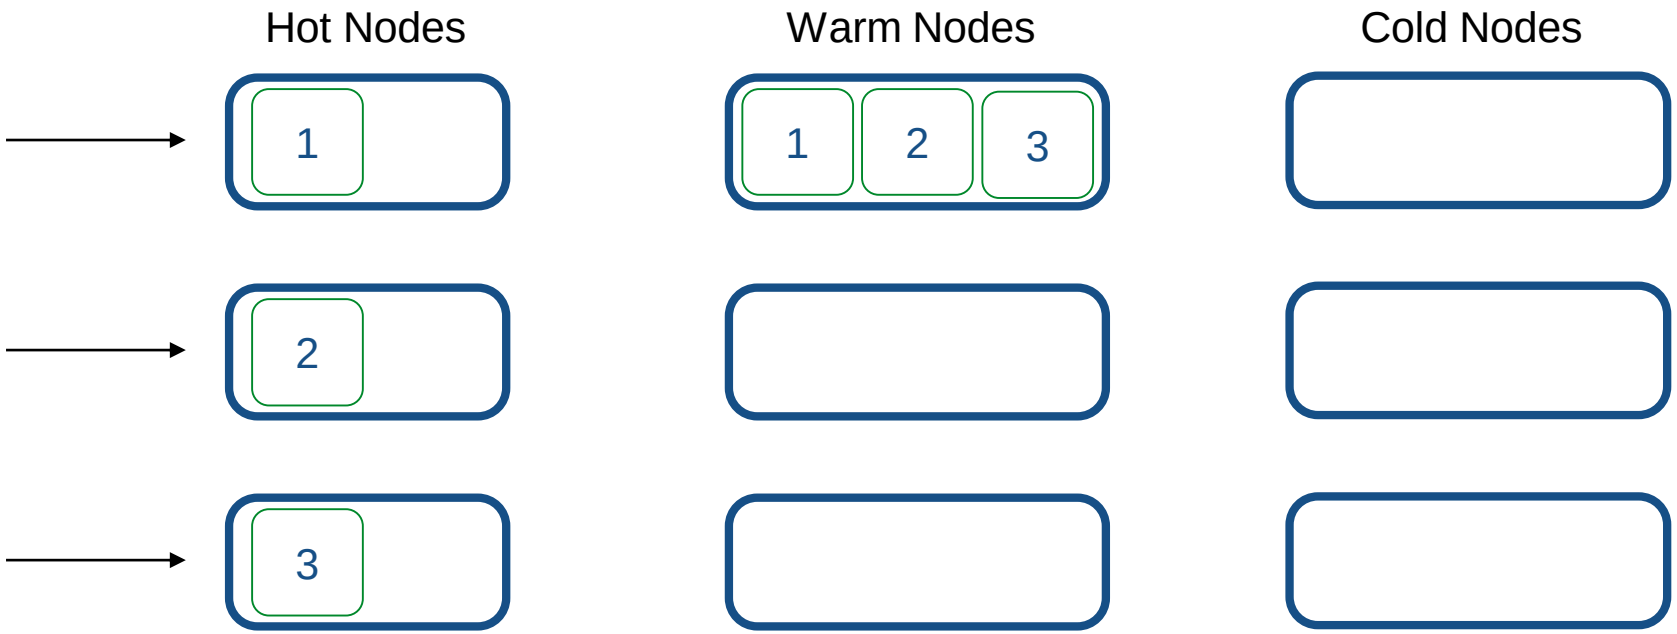

索引生命周期管理 Index Lifecycle Management

Warm Phase - Shrink

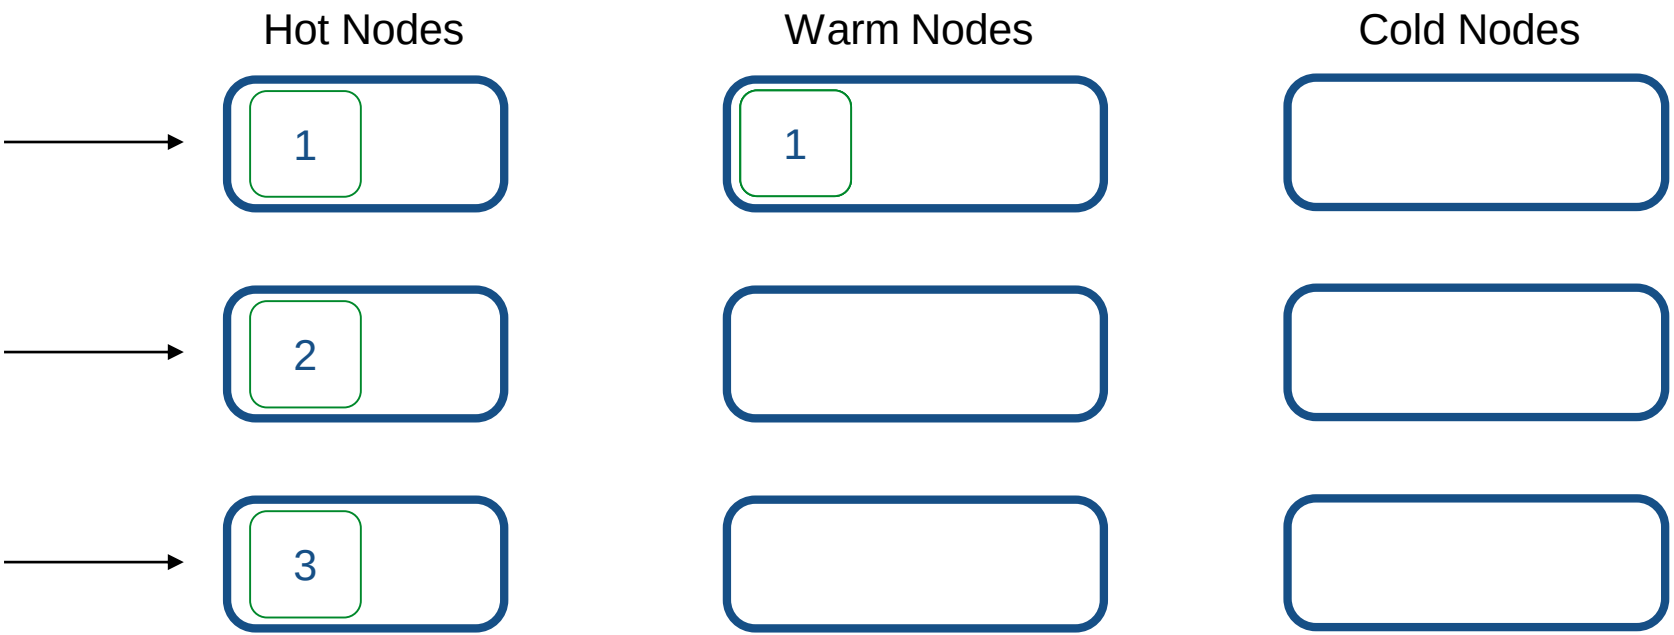

elastic

索引生命周期管理 Index Lifecycle Management

Warm Phase - Compress

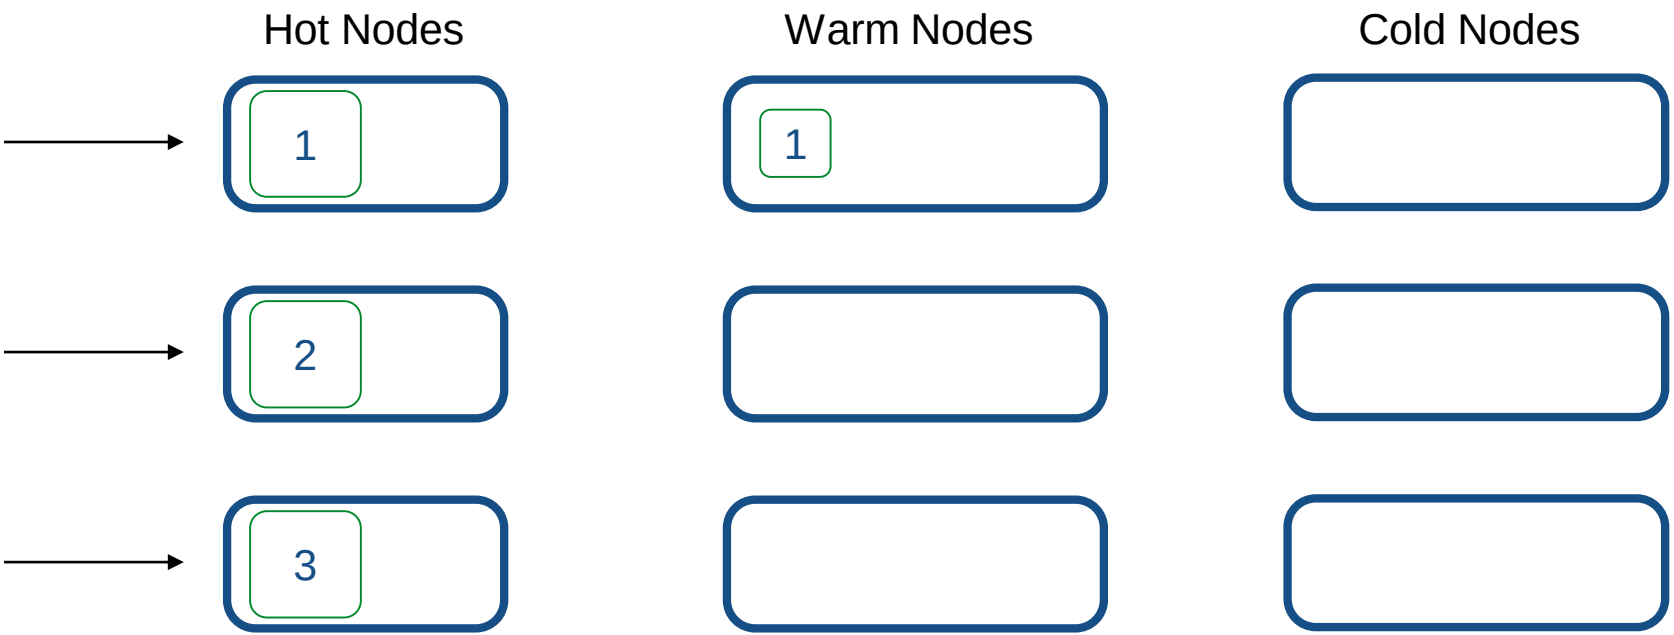

elastic

索引生命周期管理 Index Lifecycle Management

Cold Phase - Allocate

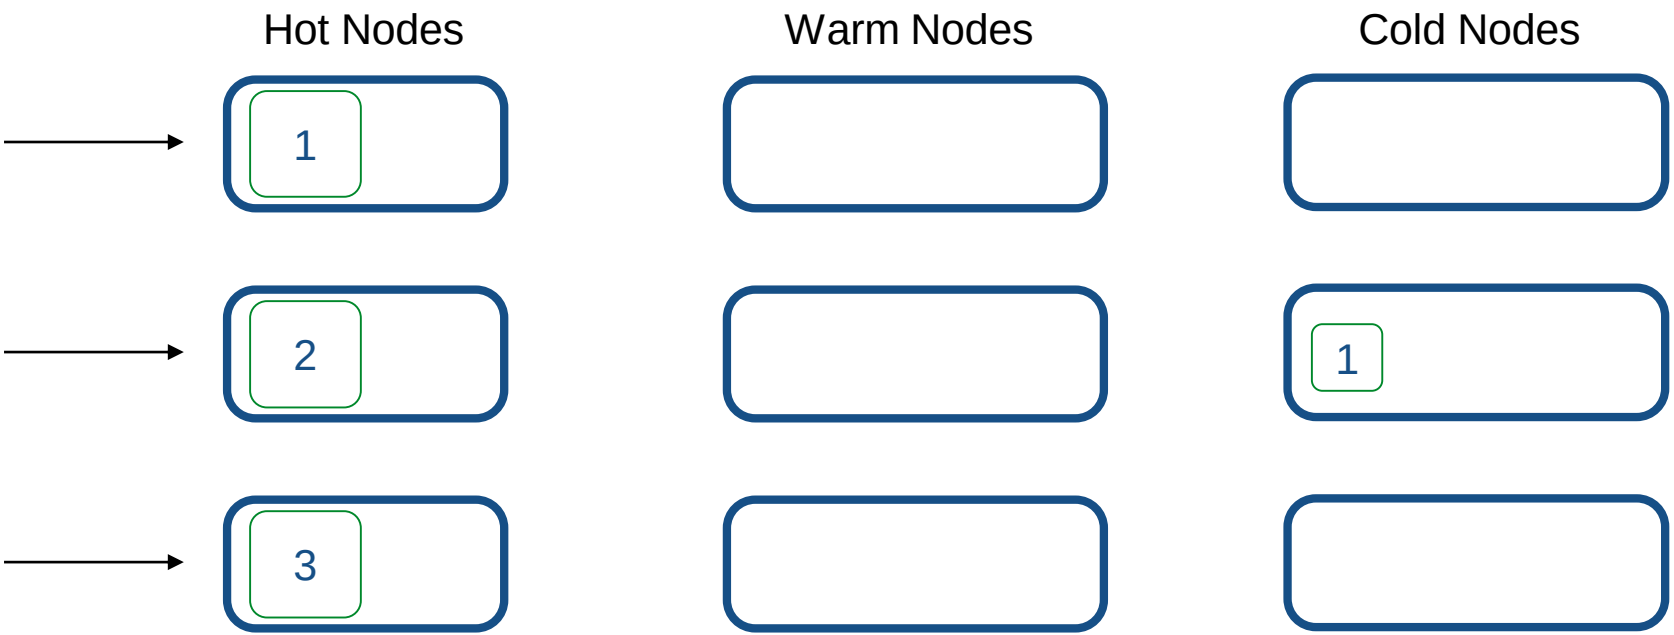

索引生命周期管理 Index Lifecycle Management

Cold Phase - Freeze

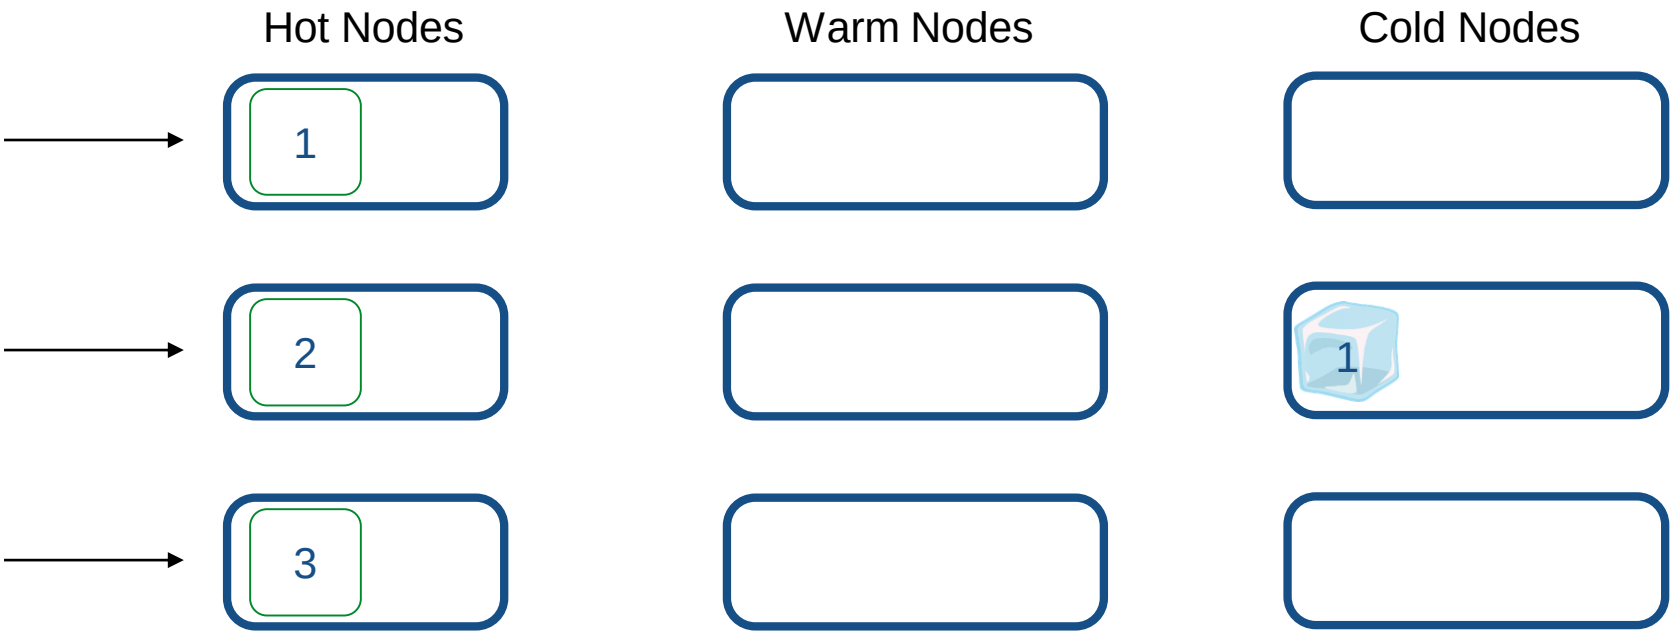

elastic

索引生命周期管理 Index Lifecycle Management

Delete Phase

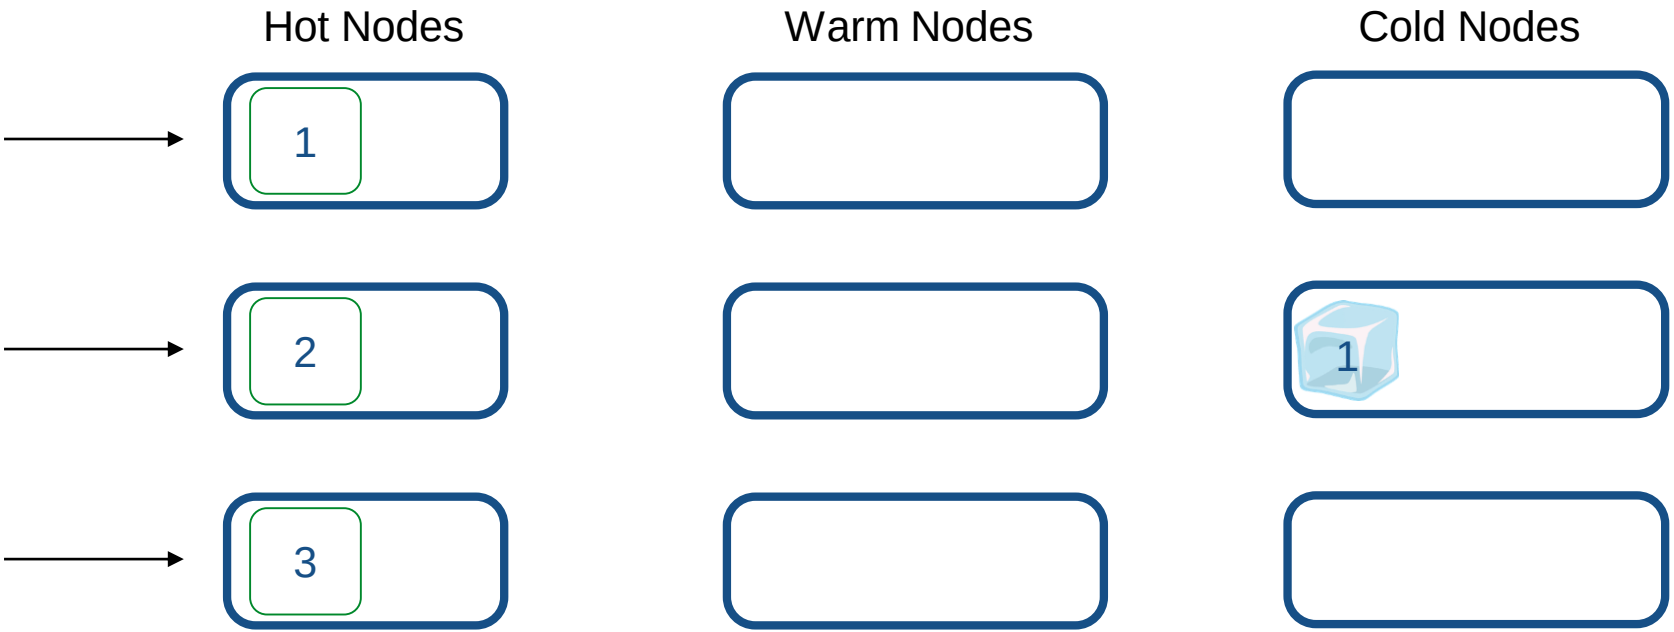

索引生命周期管理 Index Lifecycle Management

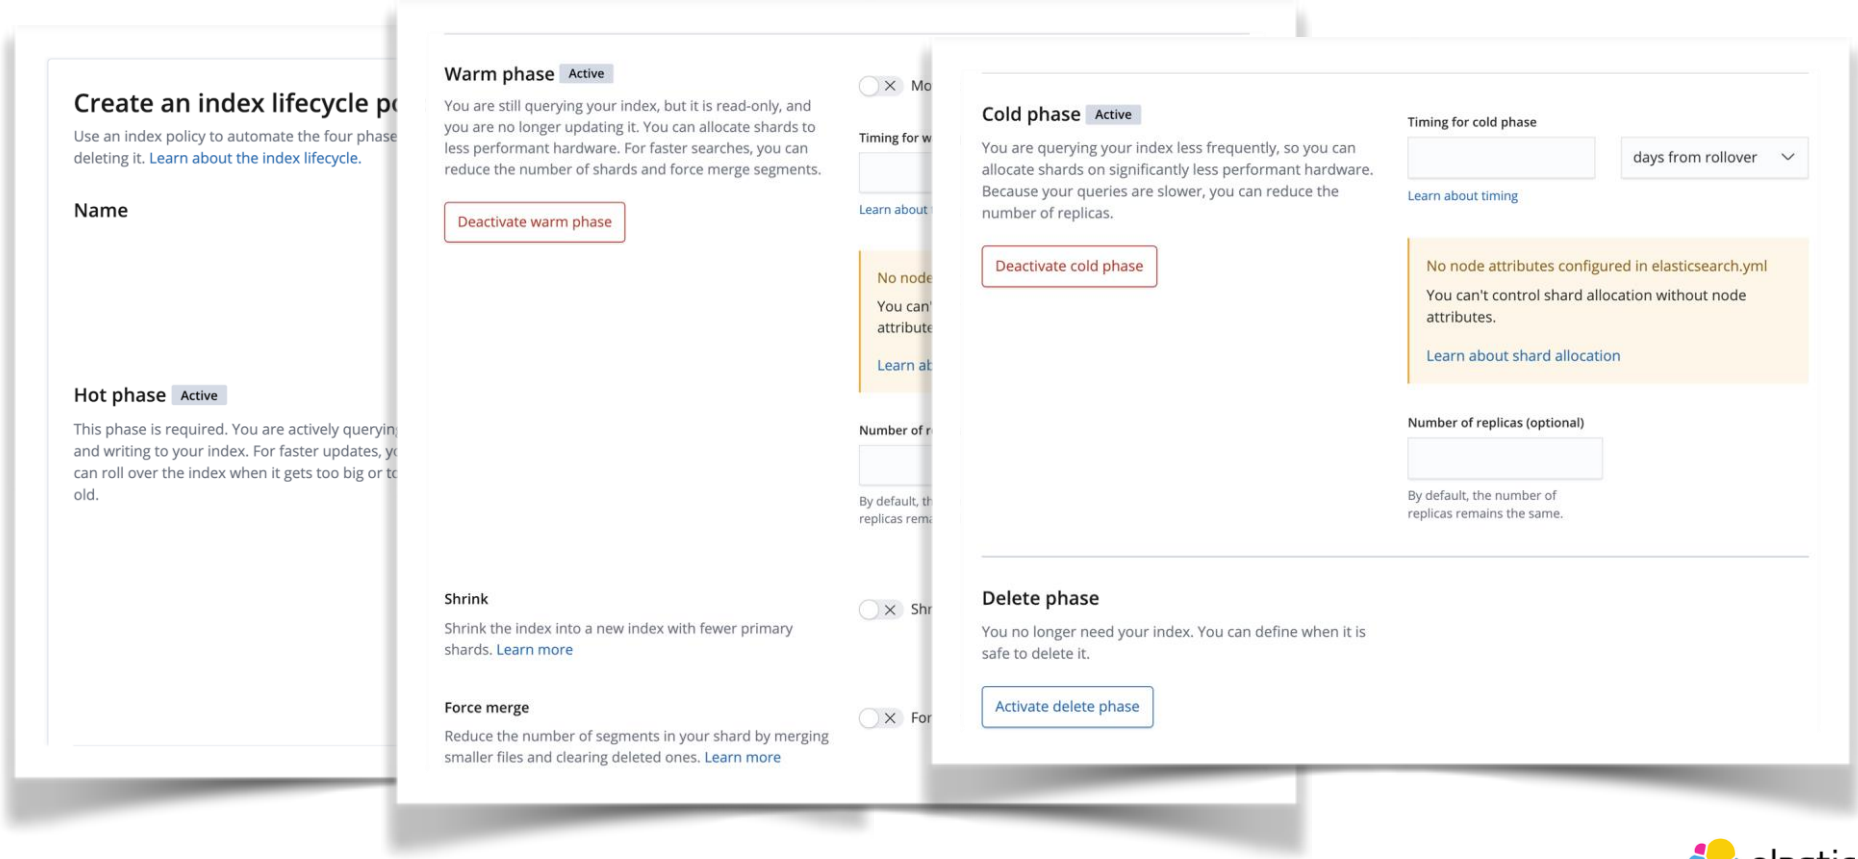

Elastic通用模式 Elastic Common Schema

好处

- 提升来自不同源头的数据关联度
- 复用分析的能力
- 复用内置可视化、仪表盘等分析能力

状态

- 当前版本1.0.0: github.com/elastic/ecs
- 内部验证
- 欢迎社区的反馈

ECS Revision: 0.992 Group 1 Fields: 3 Group 2 Fields: 81

Group 1 (Must be populated)

@timestamp
ecs_version
message

Group 2 (Must be populated to the max extent practical where event message contains relevant fields.)

Event	Device	Host	Agent	Network	Source	Destination	Service	Resource
event.category	device.mac	host.mac	agent.id	network.protocol	source.mac	destination.mac	service.id	resource.type
event.type	device.timezone_offset	host.timezone_offset	agent.name	network.forwarded_ip	source.ip	destination.ip	service.name	resource.id
event.data_source_id	device.ip	host.ip	agent.version	network.inbound_bytes	source.hostname	destination.hostname	service.version	resource.file_name
event.module	device.network_interface	host.network_interface		network.inbound_packets	source.domain	destination.domain	service.type	resource.uri
event.organization_name	device.hostname	host.hostname		network.outbound_bytes	source.port	destination.subdomain	service.state	resource.url
event.organization_id	device.type	host.id		network.outbound_packets		destination.port	service.query	resource.path
event.id	device.vendor	host.type		network.total_bytes			service.response_code	resource.version
event.raw	device.product	host.sub_type		network.total_packets				resource.hash_value
event.hash	device.version	host.operating_system		network.direction				resource.hash_type
event.tags	device.serial_number	host.operating_system_version						
event.labels	device.action	host.provider						
event.duration	device.rule_set	host.availability_zone						
event.severity	device.event_id	host.region						
event.risk_score								

elastic

Rally – Elastic官方出品的开源性能基准测试工具

https://github.com/elastic/rally

Metric	Task	Value	Unit
Total indexing time		28.0997	min
Total merge time		6.84378	min
Total refresh time		3.06045	min
Total flush time		0.106517	min
Total merge throttle time		1.28193	min
Median CPU usage		471.6	%
Total Young Gen GC		16.237	s
Total Old Gen GC		1.796	s
Index size		2.60124	GB
Totally written		11.8144	GB
Heap used for segments		14.7326	MB
Heap used for doc values		0.115917	MB
Heap used for terms		13.3203	MB
Heap used for norms		0.0734253	MB
Heap used for points		0.5793	MB
Heap used for stored fields		0.643608	MB
Segment count		97	
Min Throughput	index-append	31925.2	docs/s
Median Throughput	index-append	39137.5	docs/s
Max Throughput	index-append	39633.6	docs/s
50.0th percentile latency	index-append	872.513	ms
90.0th percentile latency	index-append	1457.13	ms
99.0th percentile latency	index-append	1874.89	ms
100th percentile latency	index-append	2711.71	ms

Agenda

Use color to highlight

1. Elastic Stack 架构
2. Elasticsearch 最佳实践
3. 如何升级
4. 性能优化
5. 问与答

升级

- 滚动升级 Rolling Upgrade
- 停机升级 Full Restart Upgrade

Upgrade - 升级助手

https://www.elastic.co/products/upgrade_guide

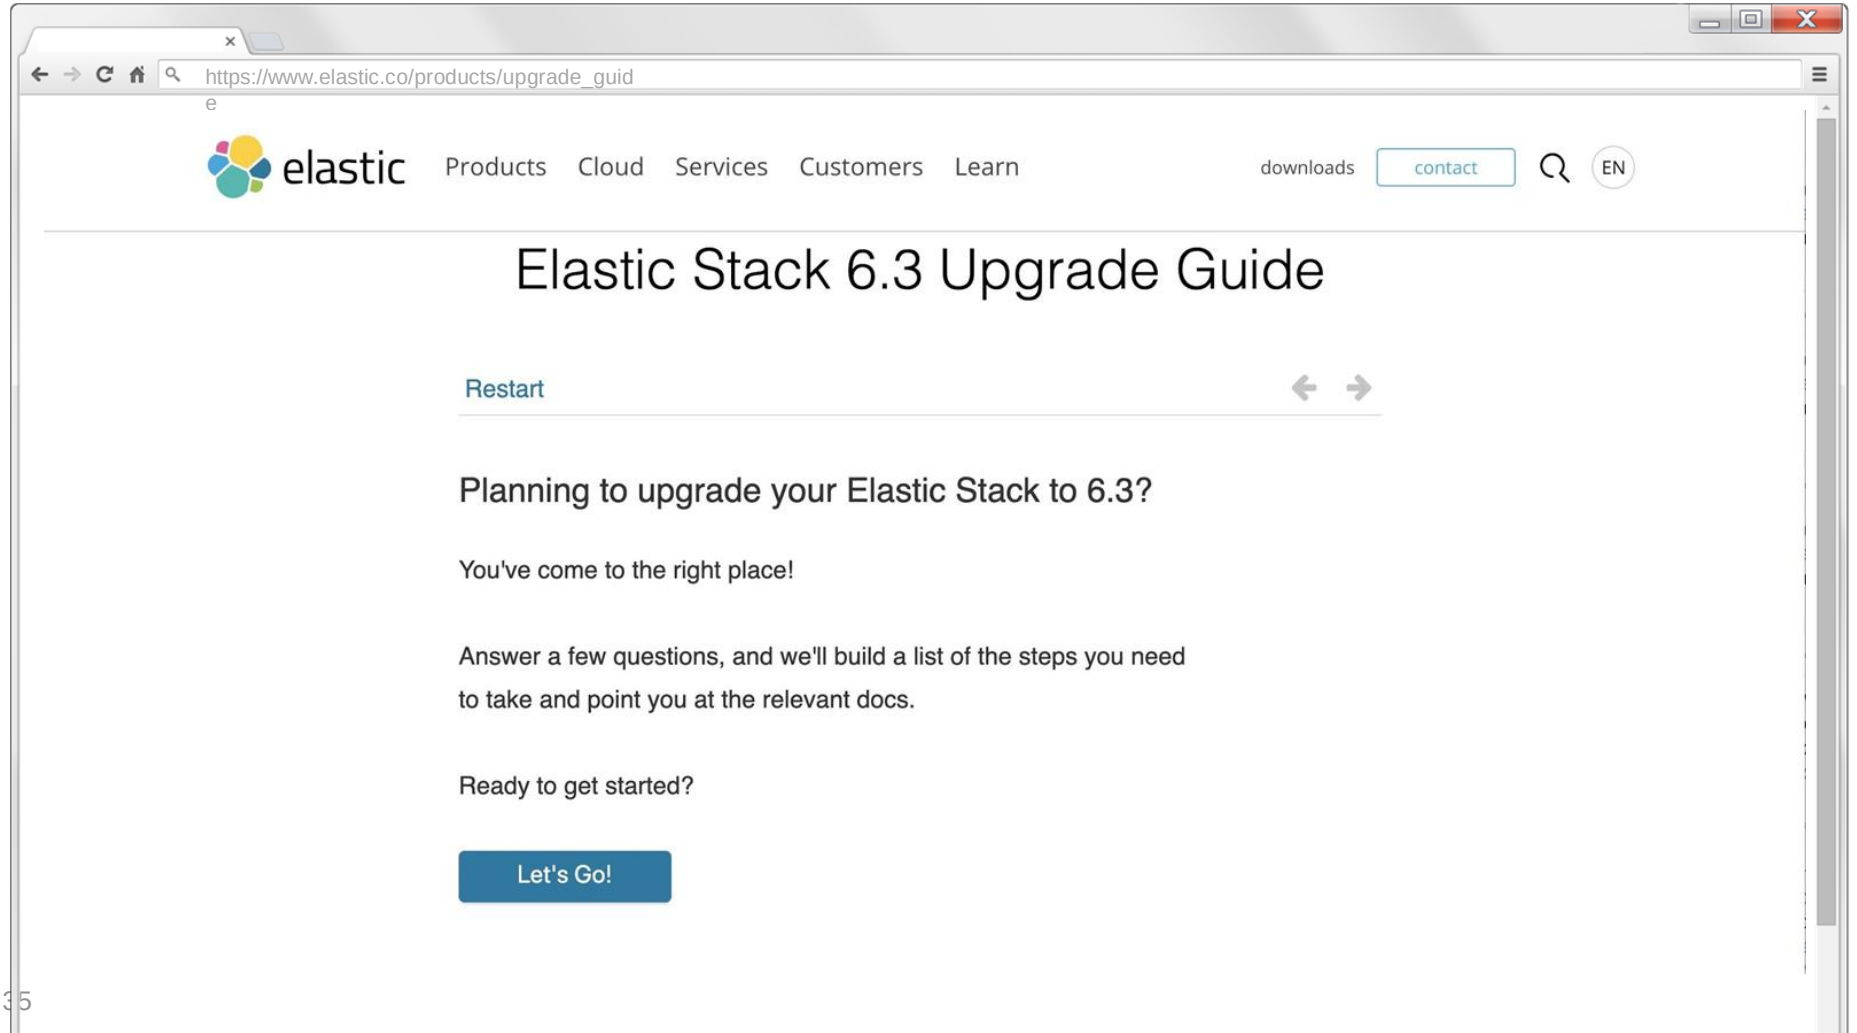
https://www.elastic.co/products/upgrade_guide
elastic
Products
Cloud
Services
Customers
Learn
downloads
contact
EN
Elastic Stack 6.3 Upgrade Guide
Restart
Planning to upgrade your Elastic Stack to 6.3?
You've come to the right place!
Answer a few questions, and we'll build a list of the steps you need to take and point you at the relevant docs.
Ready to get started?
Let's Go!

https://www.elastic.co/products/upgrade_guide

Here you go!

These are the steps you need to take to upgrade to 6.3.

Prepare to upgrade to 6.3

1. Back up your data.
2. Address any 6.0 breaking changes that affect your applications:
   - Elasticsearch breaking changes
   - Beats breaking changes
   - Logstash breaking changes
   - Kibana breaking changes
3. Check the Elasticsearch deprecation log.

Perform a Full Cluster Restart Upgrade to 6.3

4. Stop sending data to your cluster.
5. Shut down your cluster and install Elasticsearch 6.3 on all nodes.
6. Restart your Elasticsearch cluster.
7. Upgrade the internal .kibana index.
8. Upgrade Kibana to 6.3.
9. Upgrade Logstash to 6.3.

Agenda

Use color to highlight

1. Elastic Stack 架构
2. Elasticsearch 最佳实践
3. 如何升级
4. 性能优化
5. 问与答

性能优化-群集监控

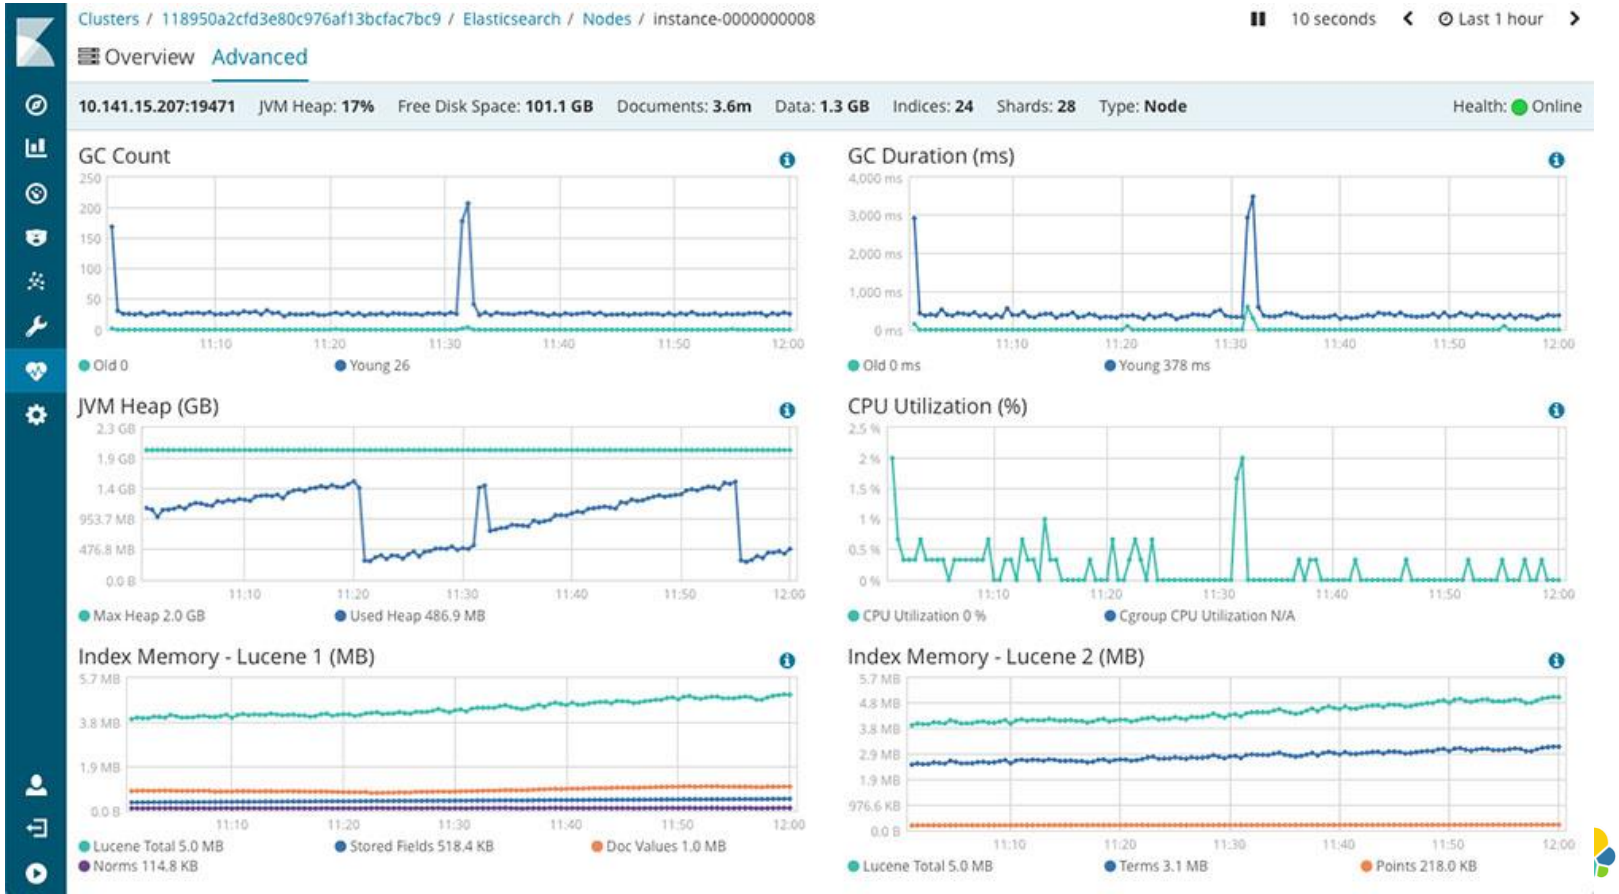

elastic

性能优化-节点监控

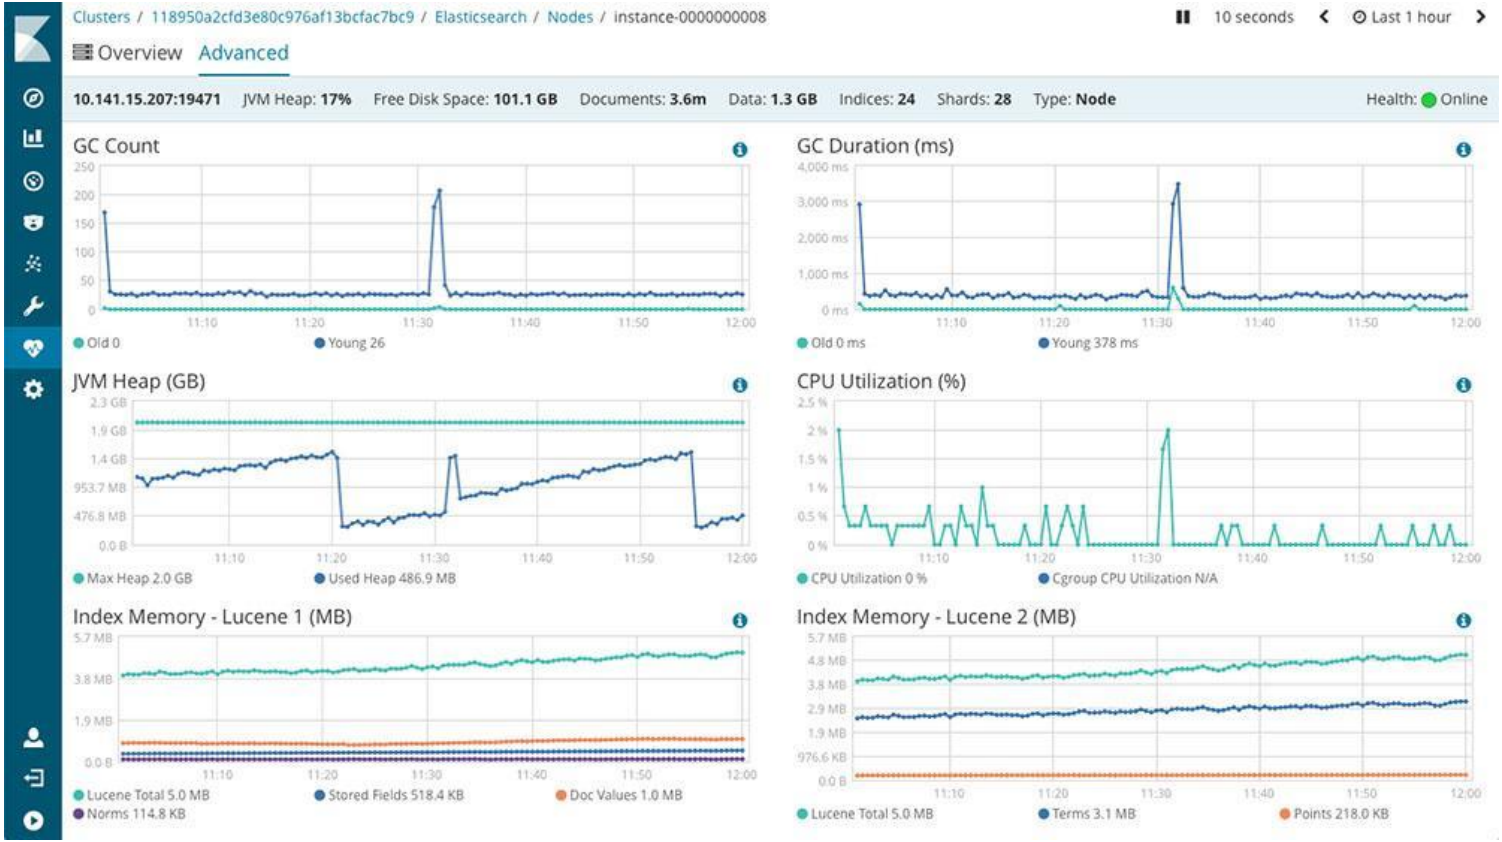
Clusters / 118950a2cfd3e80c976af13bcfac7bc9 / Elasticsearch / Nodes / instance-0000000008
10 seconds
Last 1 hour
Overview
Advanced
10.141.15.207:19471 JVM Heap: 17% Free Disk Space: 101.1 GB Documents: 3.6m Data: 1.3 GB Indices: 24 Shards: 28 Type: Node
Health: Online
GC Count
Old 0
Young 26
GC Duration (ms)
Old 0 ms
Young 378 ms
JVM Heap (GB)
Max Heap 2.0 GB
Used Heap 486.9 MB
CPU Utilization (%)
CPU Utilization 0 %
Cgroup CPU Utilization N/A
Index Memory - Lucene 1 (MB)
Lucene Total 5.0 MB
Stored Fields 518.4 KB
Doc Values 1.0 MB
Norms 114.8 KB
Index Memory - Lucene 2 (MB)
Lucene Total 5.0 MB
Terms 3.1 MB
Points 218.0 KB

性能优化-索引监控

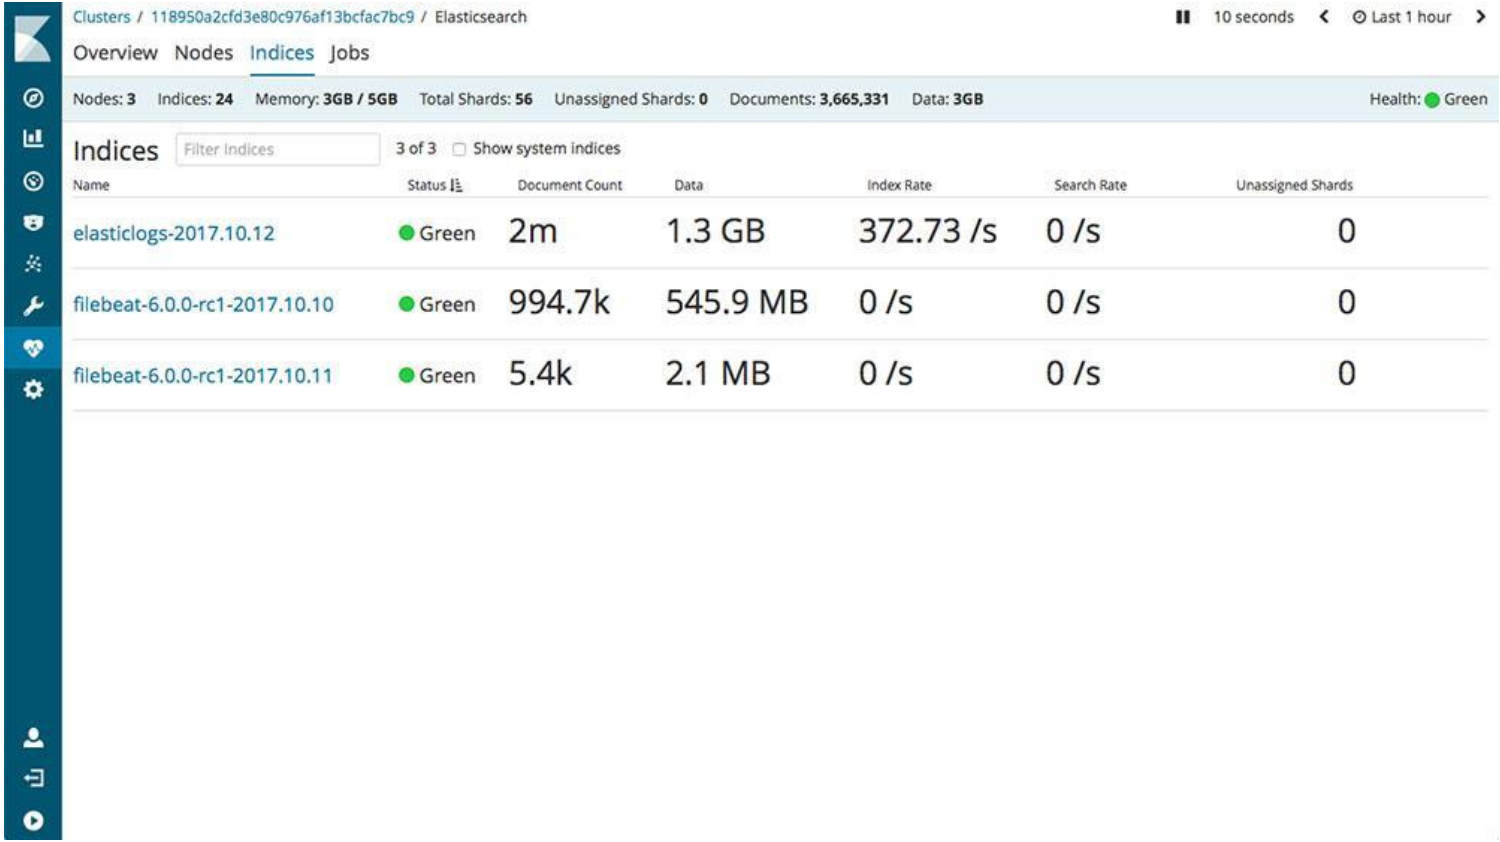

elastic

性能优化-Kibana监控

elastic

性能优化-Logstash监控

elastic

写入优化：Tune for Index Speed

https://www.elastic.co/guide/en/elasticsearch/reference/6.6/tune-for-indexing-speed.html

- 使用批量处理 Bulk Request
- index.refresh_interval 调大到 30s 或更大
- index.translog.durability:"async" index.translog.sync_interval: "10s"
- 使用自动生成的id
- 使用SSD
- 映射Mapping
    - index: false
    - norms: false
    - index_options: freqs
    - dynamic template - no default setting for string
    - index.codec: best_compression

读取优化：Tune for Search Speed

https://www.elastic.co/guide/en/elasticsearch/reference/6.6/tune-for-search-speed.html

- 数据模式设计- 避免join
- 搜索尽可能少的字段
- 避免使用脚本和预索引数据
- 对于一些枚举的数字，例如http状态代码，建议使用keyword数据类型
- 搜索 rounded dates
- 仅对只读索引进行合并索引
- 预热全局序global ordinals
- 索引排序
- 使用请求路由/优先
- 使用自适应副本选择

Elastic 专业服务

咨询服务

由Elastic的专家或者合作伙伴给订阅客户提供

为架构设计、移植和整合提供顾问服务

更多信息:
www.elastic.co/services

公开培训

目标20 - 25个人

通常由2+教师教学

全球培训课程:
purchases.elastic.co/

私塾培训

在客户场地教学

15人以上价格划算; 与公开培训教学大纲一致

需求请发送至:
training@elastic.co

elastic

elastic

AMA & Thank You!

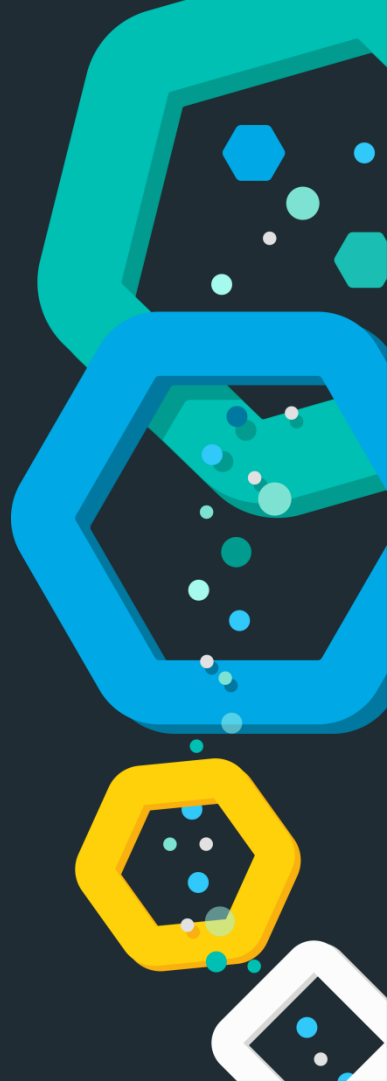

elastic
中文社区

专业、垂直、纯粹的 Elastic 开源技术交流社区
https://elasticsearch.cn/

elastic