Malicious Hubs: Detecting Abnormally Malicious Autonomous Systems

Andrew J. Kalafut
School of Informatics and Computing
Indiana University at Bloomington
akalafut@cs.indiana.edu

Craig A. Shue
Computational Sciences and Engineering
Oak Ridge National Laboratory
shueca@ornl.gov

Minaxi Gupta
School of Informatics and Computing
Indiana University at Bloomington
minaxi@cs.indiana.edu

***Abstract*—While many attacks are distributed across botnets, investigators and network operators have recently targeted malicious networks through high profile autonomous system (AS) de-peerings and network shut-downs. In this paper, we explore whether some ASes indeed are safe havens for malicious activity. We look for ISPs and ASes that exhibit disproportionately high malicious behavior using 12 popular blacklists. We find that some ASes have over 80% of their routable IP address space blacklisted and others account for large fractions of blacklisted IPs. Overall, we conclude that examining malicious activity at the AS granularity can unearth networks with lax security or those that harbor cybercrime.**

I. Introduction

The Internet is plagued by malicious activity, from spam and phishing to malware and denial-of-service (DoS) attacks. Much of it thrives on armies of compromised hosts, or *botnets*, which are scattered throughout the Internet. However, malicious activity is not necessarily evenly distributed across the Internet: some networks may employ lax security, resulting in large populations of compromised machines, while others may tightly secure their network and not have any malicious activity. Further, some networks may exist solely to engage in malicious activity. Several recent ISP enforcements, such as the Atrivo and McColo autonomous system (AS) de-peerings [1], [2] and the FTC closure of Pricewert networks [3], highlight that there are networks that exist simply to launch attacks. In this paper, we examine whether we can find malicious networks in a systematic manner using existing blacklists.

A systematic detection of disproportionately malicious networks can be used to build metrics offering several practical benefits. As an example, provider ISPs may require their customers to limit the amount of malicious activity in their networks to avoid harboring criminals. ISPs could also use the metrics to determine the effectiveness of their efforts to combat abuse and compare themselves with other networks. Also, when receiving traffic, a destination network could prioritize traffic based on the cleanliness of ASes. This would allow a network under attack to prioritize traffic that is less likely to be associated with attackers. Finally, such metrics could also aid spam filtering programs in their scoring of email messages.

To determine which ASes are malicious, we use 12 of the most commonly-used blacklists for spam, phishing, malware and botnet activities for a period of a month. These blacklists contain host names or IP addresses to be blacklisted. For host name-based blacklists, we first determine the IP addresses for each blocked host. We then use BGP routing tables to group these IP addresses into their originating ASes. Upon grouping these addresses by AS, we compare ASes by the percent of infected machines and the rate at which they are cleaned up. The key findings of our study are:

- Many ASes have a large fraction of their IP address range engaged in malicious behaviors: Two ISPs from Ukraine, one from Iran, and one from Belarus have over 80% of their IP addresses blacklisted. This raises red flags regarding their existence.
- Many ASes account for significant fractions of blacklists: Four ASes, three of which are US-based hosting providers, account for over 6% of at least one of the blacklists we tested.
- Many providers either harbor malicious activities or fail to consider them when peering: We find 22 providers with 100% of their customer ASes engaged in significant malicious activity.

Overall, these results confirm that examining malicious activity at the AS granularity can find networks with lax security or ones that harbor cybercrime.

II. Data Collection

To create a comprehensive evaluation of ASes, we use a diverse set of data sources. Each of our data sources list machines reported as engaging in some form of malicious activity. Before we describe the data sets themselves, we note their limitations: some data sets may list many IP addresses for the same compromised machine because of DHCP effects while others may group multiple compromised machines under the same address due to NAT. While important considerations, we note that these concerns are common across all networks and our analysis compares equivalently sized networks. Accordingly, while these unavoidable effects are present, they should not significantly affect our analysis.

For each data set, the data was collected from June 1, 2009 to June 30, 2009 unless otherwise indicated. We summarize

TABLE I
OVERVIEW OF DATA SETS

Label	Description	Duration (days)	Unique IP Addresses	Unique ASes
APWG	Phishing URLs from the Anti-Phishing Working Group	30	9,560	1,803
Bot C&C	Botnet command and control IPs from the ShadowServer Foundation	30	1,986	611
CleanMX	Malware serving sites from the CleanMX VirusWatch mailing list	30	2,974	687
eSoft	Malware serving sites from eSoft, Inc.	30	8,000	1,196
Local Spam	URLs from spam messages received by the IU CS Department	30	5,495	1,024
Malware Patrol	MalwarePatrol's block list for malware-serving sites	30	871	368
PhishTank	Phishing URLs from PhishTank	28	7,143	1,580
Spamhaus SBL	Verified spam sources from Spamhaus.org Block List	29	6,422	2,005
Spamhaus XBL	Hijacked machines from Spamhaus.org Exploit Block List	29	29,585,604	13,580
SI-Feed	URLs and IP addresses from spam emails from Support Intelligence	30	7,591	1,420
SI-DNS	IP addresses from DNS resolutions on the `SI-Feed` data set	30	4,448	911
SURBL	Host names appearing in spam messages from SURBL	30	29,324	2,739

these data sources in Table I, and describe them below.

1) Phishing Sites: Phishing web sites attempt to collect sensitive data, such as login credentials, from users by impersonating legitimate organizations. The Anti-Phishing Working Group (APWG) [4] and PhishTank [5] have among the largest data feeds listing such phishing sites. We have access to this data and use it to create our `APWG` and `PhishTank` data sets, respectively. Both of these feeds contain URLs of phishing sites, along with other metadata. On an hourly basis, we extract host names from the URLs currently in the feed, and perform DNS resolutions in each host name to get lists of IP addresses associated with these feeds. The PhishTank data set had a two-day outage on June 20 and June 21 causing us to only have 28 days of data for that data set.

2) Spam/Scam Sites: Similar to their phishing site brethren, scam sites are often advertised in unsolicited messages. These spam-advertised sites may actually be phishing sites, be involved in some other type of scam, or provide actual legitimate products or services. Two of the major providers of lists of such sites, Support Intelligence [6] and SURBL [7], have granted us access to them.

We receive the feed from Support Intelligence every six hours. This feed contains URLs from spam as well as associated IP addresses. We use these IP addresses as our `SI-Feed` data set. Not every URL in this feed has an associated IP address, and for some that do, when we resolve the associated host names we get different addresses. Therefore, we use our own resolutions of these as another data set, `SI-DNS`.

SURBL also collects domain names from URLs contained in spam. Although they typically only allow users to perform look-ups on the domain names in their list, we have also arranged to receive the associated IP addresses from them. These IP addresses are those associated with the domain itself, and with the domain with `www` prepended. We receive this feed once per day, and refer to it as `SURBL`.

Finally, we harvest URLs from spam sent to the Computer Science at Indiana University (IU) and use it to create the `Local Spam` data set. We receive the list of URLs appearing in spam on a daily basis and extract the host names and perform DNS resolutions to obtain the IP addresses.

3) Spam Senders: A popular anti-spam approach, IP blacklisting, is often used at mail servers to prevent compromised machines from sending mail directly. Spamhaus runs the most widely-used blacklist in this context, the SBL [8]. The SBL contains IP addresses of machines verified as spam senders. This list can be queried by mail servers when they receive connections to block known spammers. We obtain a copy of this blacklist every hour, and extract the IP addresses to create the `Spamhaus SBL` data set. Data collection for the `Spamhaus SBL` data set started a day later than the others, beginning on June 2, 2009.

4) Exploited Hosts: Spamhaus also maintains a second blacklist, known as the XBL [9]. This list contains IP prefixes (often individual IP addresses) of hosts infected with exploits often used to send spam. This includes open proxies, computers infected with viruses that are known to send spam, and other exploits. This data is updated every half hour, and is labeled `Spamhaus XBL`. Data collection for this data set started a day later than the others, beginning on June 2, 2009.

5) Malware Downloads: Malicious software, or *malware*, including viruses, worms, and trojans, have harmful effects on the computers they infect. Three of our data sets list web sites which host malware downloads.

The Clean-MX Viruswatch mailing list [10], eSoft [11], and Malware Patrol [12], independently collect URLs which host malware. The Viruswatch mailing list periodically sends out emails indicating newly discovered URLs with viruses. We receive mails from eSoft with new URLs containing malware, along with a malware sample, as they are discovered. We download new URLs from Malware Patrol every hour. In each case, we extract host names and perform DNS resolutions to obtain the set of IP addresses we use. We label these data sets `CleanMX`, `eSoft`, and `Malware Patrol`, respectively.

6) Bot Command and Control: Botnets consist of groups of compromised machines used for malicious purposes on the Internet. Miscreants often use them for sending spam and for hosting phishing and scam sites. While we do not have any direct sources of botnet IP addresses, many of the addresses in our other data sources are likely to be bots since bots are commonly used for malicious activity. However, botnets must get their instructions from their bot masters, often through command and control servers, which distribute orders. The ShadowServer Foundation [13] provides lists of botnet command and control servers along with their IP addresses. We have access to this data and update it hourly. We refer to this data set as `Bot C&C`.

III. Degree of Autonomous System Maliciousness

From the IP addresses from our data sets, we can determine the originating AS for each, and use this to group hosts at the AS granularity. In order to map IP addresses to an AS, we used a June 15, 2009 BGP routing table from the RouteViews Project [14]. We chose this date because it is in the middle of our data collection and is expected to give us the best estimate of the routing information from that duration. We loaded each advertised BGP prefix and originating AS from the RouteViews data into a trie data structure commonly used by the routers in deciding the next interface to use to forward packets. We then performed longest prefix matches on each IP address to determine the AS associated with the address. Using the AS information corresponding to each malicious IP, we examined the extent of AS maliciousness from two perspectives: the percentage of IP address space for an AS found to be blacklisted and the percentage of the blacklist each AS constitutes. We now describe both approaches and their results in detail.

A. Examination of ASes by Fraction of Advertised IP Space

Given the number of malicious IP addresses associated with an AS, the most straight-forward approach to evaluating the ASes for maliciousness would be to simply order the ASes by number of malicious IP addresses. However, such an analysis would penalize the larger ASes: they simply have more addresses so they have more hosts that could be compromised and blacklisted. Accordingly we must consider the overall size of the AS as a factor when looking for ASes that are disproportionately bad.

There are no direct sources that help estimate the size of an AS. However, the prefixes advertised by an AS can be used to determine the maximum number of routable IP addresses associated with the AS. While ASes often have unused IP addresses in each of their prefixes, and it is difficult to determine just how many addresses are unused, this allows us to obtain a rough upper bound for the AS size. We again extracted the prefix and originating AS information from the June 15, 2009 BGP RouteViews routing table. We loaded this information into a trie data structure as before. For each prefix associated with an originating AS, this allowed us to determine the number of IP addresses associated with the prefix. In the process, we were careful to exclude any sub-prefixes associated with other ASes. After adding together the address space from each of the prefixes for each AS, we had the total number of IP addresses advertised by each AS.

With our information about the number of unique machines found in at least one of our data sets and the rough size of each AS, we can determine the rough percentage of each AS that appears in each data set. In Figure 1, we show the percentage of badness for each AS present in our data sets, excluding the Spamhaus XBL data set. We separated out the Spamhaus XBL due to its much larger size which made the other results difficult to read. This Figure shows several interesting results. First, a total of 31,263 ASes were advertised in our BGP routing data and 46.8% of these had at least one malicious IP in them. *While a majority of them have little to no malicious activity, a small number of ASes have as much as 0.5-10% of their IPs engaged in maliciousness.* In fact, in the SI-Feed data set, one AS had 9.25% of its addresses in the data set. No other AS had 5% or more of its addresses in any of these data sets.

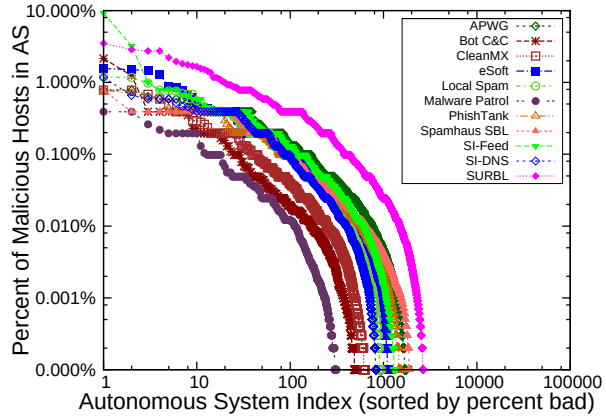

Fig. 1. Percentage of badness for each AS. The AS indices are sorted from the most malicious AS to the least malicious for each data set.

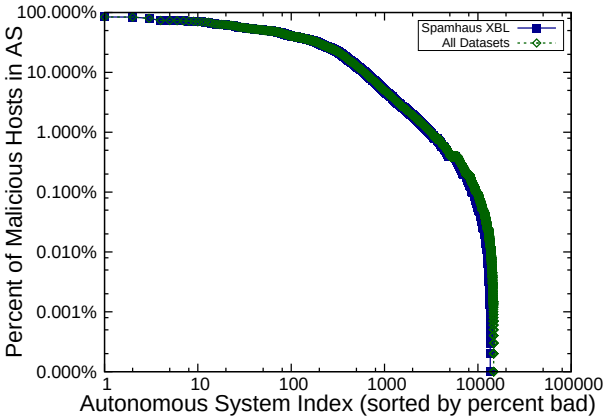

Fig. 2. Percentage of badness for each AS in the Spamhaus XBL blacklist and across all blacklists combined.

In Figure 2, we show the same results for the Spamhaus XBL data set and the combination of each data set. We note that the two lines are very similar and almost completely overlap because of the size of the Spamhaus XBL data set. We found 58 ASes with over 100,000 compromised machines in this data set. Additionally, 255 ASes had between 10,000 and 100,000 machines blacklisted. When looking at the percentage of each AS's advertised address space marked as malicious, we found that *four ISPs, two from Ukraine, one from Iran, and one from Belarus, had at least 80% of their advertised IP space blacklisted*. Another 49 in the Spamhaus XBL data set had 50-80% of their addresses listed. Further, 556 ASes had at least 10% but less than 50% of their IP addresses listed. This indicates that some ASes have either too lax a security policy or may be intentionally harboring cybercrime.

TABLE II
NUMBER OF ASES IN EACH DATA SET CONTAINING THE GIVEN PERCENTAGE OF ALL IP ADDRESSES IN THE DATA SET.

Percent of IPs	APWG	Bot C&C	CleanMX	eSoft	Local Spam	Malware Patrol	PhishTank	Spamhaus SBL	Spamhaus XBL	SI-Feed	SI-DNS	SURBL
$\geq$ 10%												
[9%, 10%)		1										
[8%, 9%)		1										
[7%, 8%)									1			
[6%, 7%)										1		
[5%, 6%)				1			1					
[4%, 5%)	1	1	2						1	1		1
[3%, 4%)					3	1			1		1	2
[2%, 3%)	2	2	2	3	2	1	1		3	1	2	
[1%, 2%)	5	5	3	7	11	6	3		7	5	10	8
[0.50%, 1%)	12	10	16	6	19	16	11		16	20	19	14
[0.25%, 0.50%)	20	26	27	25	20	18	18	18	18	27	33	38

B. Examination of ASes by Proportion of Data Set

While examining the percentage of an AS that is blacklisted can highlight ASes with disproportionately high concentrations of blacklisted hosts, it requires large data sets. While the Spamhaus XBL data set shows this clearly, other data sets are not large enough to distinguish atypically malicious networks. However, rather than consider the AS to be malicious based on the percentage of its blacklisted address space, we can instead examine the percentage of a data set that an AS represents. This can be used to highlight ASes with a large number of blacklisted hosts.

We begin by finding the number of ASes containing at least 0.25% of the IP addresses in each data set. These results are shown in Table II. In doing so, we wanted to avoid penalizing large ASes that advertise large address spaces and do not necessarily account for a disproportionate amount of maliciousness in that data set. Toward that goal, we first find the percent of data set belonging to each AS. Then we find the fraction of IP address space this AS has with respect to all ASes represented in the data set. If the first is a factor of 10 greater than the second, we take the AS into account. Otherwise, we ignore it. For example, if an AS contained exactly 0.25% of the IPs in the data set, we would list it if it accounted for less than 0.025% of the address space of all ASes in the data set, but ignore it otherwise.

We see from the table that some ASes have high concentrations of malicious activity. *For example, in the Bot C&C data set, we see that one AS contains 9.11% of the IP addresses in the data set, yet its advertised address space represents only 0.002% of the address space advertised by all ASes in the data set.* The next AS in this data set, with 8.66% of the listed IP addresses represents only 0.006% of the advertised addresses in the listed ASes. Of these two ASes, one is a large broadband ISP from Turkey and the other is a hosting service provider from the US. Incidentally, the US-based hosting provider also accounts for 7-8% of all blacklisted IPs. Further, in Spamhaus XBL and SI-Feed data sets, we find two more US-based hosting providers that account for over 6-8% of these blacklists.

Overall, a small number of ASes have a disproportionate fraction of malicious hosts. These ASes may harbor malicious activity and should be investigated similarly to Atrivo or McColo [1], [2]. We believe that legitimate ISPs with disproportionately high malicious activity need to provide tighter account controls, or seek opportunities to provide anti-virus or firewalling services to prevent malicious activity.

C. ASes with Unruly Children

Our data establishes that malicious activity is often disproportionately clustered at a small number of ASes. We now look at whether ASes with disproportionate malicious activity are tightly clustered. We begin by labeling as malicious any AS with at least 1% of its IP addresses appearing in any blacklist. We then examine each of the BGP updates for June 2009 to find provider-customer (or parent-child) relationships. For each provider AS, we consider the extent to which its customer ASes have been found to be malicious. In the second column of Table III, we show the number of provider ASes with at least three children that have the indicated percentage of its children as malicious. *We see 22 ASes with 100% of their customers classified as malicious. A total of 194 providers have at least 50% malicious customer ASes*.

TABLE III
PERCENTAGE OF MALICIOUS CUSTOMER ASES FOR PROVIDERS WITH MORE THAN THREE CUSTOMERS.

Percent of Malicious Customer ASes	Number of Provider ASes	
	Fraction of Advertised IP Space	Proportion of Data Set
100%	22	
[90%, 100%)	2	
[80%, 90%)	8	
[70%, 80%)	17	
[60%, 70%)	72	3
[50%, 60%)	73	2
[40%, 50%)	78	5
[30%, 40%)	202	24
[20%, 30%)	239	45
[10%, 20%)	204	78

We repeated this analysis using the definition of maliciousness from Section III-B: the AS must have at least 0.25% of the malicious IPs in a data set. We show these results in the third column of Table III. *Five providers have at least 50% of their customer ASes labeled as malicious*.

This analysis shows that there are dense clusters of malicious activity in the Internet. This may be an indication that there are upstream providers that are willing to peer with any customer, regardless of whether it harbors malicious activity. We hope that studies similar to ours would put pressure on provider ASes to extensively screen their customers and require their customers to limit malicious activity as part of their peering agreements.

IV. Related Work

Some previous works attempt to locate malicious behavior at granularities other than ASes. In their study of spyware, Moshchuk *et al.* [15] find that certain categories of web sites contain more spyware than others. Similarly, work by Provos *et al.* [16] finds that 67% of malware download sites in drive-by downloads are hosted in a single country, China. While there is insight to be gained by examination at these other granularities, we focus solely on the AS location of malicious behavior in the paper.

Other work touches on AS locations of malicious behaviors on the Internet. In a paper on spammers' behaviors, Ramachandran *et al.* [17] find that a small number of ASes are responsible for sending a large amount of spam, with 36% of all spam coming from just 20 ASes. Konte *et al.* [18] examined scam hosting infrastructure. Among their findings was that for the spam campaigns they examined there was almost no overlap in the ASes of the spamming machines and the ASes where the scam web sites were hosted. However, none of these papers has the AS locations of the behavior as their main focus as we do.

Numerous studies have focused on accurately determining types of AS relationships, including those by Di Battista *et al.* [19], Dimitropoulos *et al.* [20], Gao [21], and Subramanian *et al.* [22]. Where we deal with connections between ASes, we are most concerned just with if a malicious AS is related to other malicious ones. Therefore to infer the type of relationship, we use a simple algorithm similar to the one Gao describes as her basic algorithm.

V. Conclusion

In this preliminary work, we examined whether some networks serve as safe harbors for malicious activity. We found that several ASes have high concentrations of malicious IPs while others represent disproportionately higher malicious activity than their equivalently sized peers. This shows that while botnets are commonly being used to launch attacks, malicious hosts may still be clumped by network providers. In spite of these results, traffic cannot simply be declared malicious based solely on its originating AS even for ASes with the high degree of maliciousness, as this would have extensive collateral damage, penalizing legitimate traffic as well. However, identifying if traffic is coming from ASes known to be malicious can be used as one component to help make such a decision.

There are several interesting open questions about malicious ASes which we plan to address in future work. First, we took two approaches towards identifying malicious ASes. Other approaches are possible and should be explored. Additionally, we plan on examining other characteristics of malicious ASes such as their BGP behaviors. A more in-depth analysis to be able to understand the motivation behind these AS behaviors. It will also help differentiate ones that actually belong to miscreants from those that just ignore malicious activity. We expect that our analysis to increase ISP accountability. It can become part of a mechanism to combats malicious activity. By providing a comparison with equivalently-sized networks, we can highlight ASes in most need of attention. This information can also be used in peering agreements to place pressure on ISPs to respond to malicious activity.

Acknowledgments

We would like to thank the RouteViews project for their extensive publicly available BGP data. We would also like to thank the providers of all the lists of malicious IP addresses and URLs we used.

References

[1] J. Hruska, "Bad seed ISP Atrivo cut off from rest of the Internet," 2008. [Online]. Available: http://arstechnica.com/security/news/2008/09/bad-seed-isp-atrivo-cut-of%f-from-rest-of-the-internet.ars
[2] ——, "Spam sees big nosedive as rogue ISP McColo knocked offline," 2008. [Online]. Available: http://arstechnica.com/security/news/2008/11/spam-sees-big-nosedive-as-%rogue-isp-mccolo-knocked-offline.ars
[3] J. Cheng, "FTC forces hive of scum and villainy ISP offline," 2009. [Online]. Available: http://arstechnica.com/tech-policy/news/2009/06/ftc-forces-hive-of-scum%-and-villainy-isp-offline.ars
[4] APWG, "Anti-phishing working group," http://www.antiphishing.org/.
[5] OpenDNS, "PhishTank," http://www.phishtank.com/.
[6] Support Intelligence, LLC, http://www.support-intelligence.com/.
[7] SURBL, http://www.surbl.org/.
[8] Spamhaus Project, "SBL," http://www.spamhaus.org/sbl/index.lasso.
[9] ——, "XBL," http://www.spamhaus.org/xbl/index.lasso.
[10] NETpilot GmbH, "Viruswatch mailing list," http://lists.clean-mx.com/cgi-bin/mailman/listinfo/viruswatch.
[11] eSoft Inc., http://www.esoft.com/.
[12] MalwarePatrol, "Malwarepatrol - malware block list," http://www.malwarepatrol.net/lists.shtml.
[13] ShadowServer Foundation, http://www.shadowserver.org/wiki/.
[14] University of Oregon Advanced Network Technology Center, "Route Views project," http://www.routeviews.org/.
[15] A. Moushchuk, T. Bragin, S. Gribble, and H. Levy, "A crawler-based study of spyware on the web," in *Internet Society Network and Distributed System Security Symposium (NDSS)*, 2006.
[16] N. Provos, P. Mavrommatis, M. Rajab, and F. Monrose, "All your iFRAMEs point to us," in *USENIX Security Symposium*, 2008.
[17] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," in *ACM SIGCOMM*, 2006.
[18] M. Konte, N. Feamster, and J. Jung, "Dynamics of online scam hosting infrastructure," in *Passive and Active Measurement Conference (PAM)*, 2009.
[19] G. D. Battista, M. Patrignani, and M. Pizzonia, "Computing the types of the relationships between autonomous systems," in *IEEE Conference on Computer Communications (INFOCOM)*, 2003.
[20] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley, "AS relationships: Inference and validation," *ACM SIGCOMM Computer Communications Review (CCR)*, vol. 37, no. 1, pp. 29–40, Jan. 2007.
[21] L. Gao, "On inferring autonomous system relationships in the internet," *IEEE/ACM Transactions Of Networking*, vol. 9, no. 6, pp. 733–745, Dec. 2001.
[22] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *IEEE Conference on Computer Communications (INFOCOM)*, 20002.